RESEARCH PAPER · NOVEMBER 2026

Verified Against What

ad verification sells buyers a verdict on whether an impression was viewable, valid, safe and suitable. This paper tests, product by product, what each verdict is accredited for, what independent evidence supports it, what third parties can see inside the large platforms, and whether a verdict changes what anyone pays.

Evgeny Popov · No Fluff Advisory

RESEARCH CUTOFF 11 NOVEMBER 2026 · ONLY INFORMATION AVAILABLE BY THAT DATE IS USED

1458	804	17
sources opened and registered	claims in the research ledger	vendors scored

20
corrections logged

ABSTRACT

Abstract

What does an ad-verification verdict establish, who can check it, and when does it decide money? This paper examines the US market through public product documents, standards, studies and contracts. Its evidence cutoff is 11 November 2026; status-sensitive facts are rechecked on 11 December 2026 and any change is published as a dated correction. It keeps measurement scope, detection evidence and contractual authority separate.

IN SHORT

The merged sold-unit grid contradicts H1a: 108 of 298 documented units hold MRC accreditation, or 36.2% ◆. H1b is contradicted on the author's readings of Peer39's pre-bid wording and Protected by Mediaocean's CTV display claim, rather than on a regulator's finding ◆. H2 is supported on the author's registered listed-terms reading, with determining effect in 2 of 16 readable entries and wider-contract sensitivities reported alongside ◆. The independent sources screened supplied no qualifying sold-category error-rate test for a scored product, supporting H3 ◆. TikTok's documented measurement-kit route contradicts H4's universal platform-supplied-data claim, at medium confidence ◆. DoubleVerify's separation statement and ISBA's review pledge contradict H5 and H6 under the author's stated interpretations, with stricter readings disclosed ◆ ◆. H7 remains unresolved because production claims and partial buyer reporting leave the verification-verdict-to-agent-decision link unestablished ◆.

The question and the method

A verification verdict can classify traffic, measure an opportunity to see an ad or apply a content policy. The buyer also needs its observation route, measurable population, selected settings and permitted action. A vendor badge alone supplies none of those fields [1, 2, 3]. The author's essays on verification's pricing power and advertising's missing adjudicator are rival framings tested against public terms, not corroborating evidence [4, 5].

The hypotheses and scoring anchors were registered on 9 November 2026. AI agents produced 10 workstreams and 7 deep profiles under the author's direction. The source register holds 1458 records; the claim ledger holds 804 claims. Claude agents researched the workstreams and profiles and scored the calibration rows with separate refute-mode verifiers. An OpenAI model completed the later sweep checks, scored and self-checked the added rows, and fact-checked the chapters. The author checked the work and made the inclusion and verdict calls. Chapter 13 distinguishes those roles from a blind second score.

What the paper found

H1a: accredited coverage is below the rule

The common product-by-metric-family-by-environment grid records 80 accredited units out of 179 sold by suites (44.7%), 25 of 82 for IVT specialists (30.5%) and 3 of 37 for suitability specialists (8.1%). Pooled coverage is 108 of 298 (36.2%). H1a requires at least 75% pooled and no cohort below 50%; the merged result fails both conditions ◆. illumaSentinel's sold-unit denominator remains unknown, rather than zero. The result describes the documented grid and MRC scope, not a spend-weighted share of impressions. Chapter 2.

H1b: scope wording requires interpretation

The author reads Peer39's reference to accredited measurement within a pre-bid signal as implying a pre-bid accreditation absent from the applicable Adloox record. A narrower reading could mean an accredited post-serve input feeding that signal. Protected by Mediaocean's CTV display-impression wording also exceeds scope on a literal reading. Those author interpretations make H1b contradicted; they establish no MRC compliance decision ◆. Read the letter for the product, metric, environment, format, placement, language and standard. Chapter 2.

H2: some reports decide money

The listed-terms test supports H2 at 2/16, or 12.5%, under the author's ruling. The Trade Desk's supply-side clause gives accredited reports conclusive effect; Microsoft Invest had a partner-report refund route whose end Microsoft announced. These counterexamples defeat a universal claim that verification moves nothing ◆. Wider public-contract populations and review-qualified or trigger-only clauses can change the share across the 25% threshold. No paid third-party-verdict settlement or matched price difference was found in the screened record; that bounded absence cannot establish what happens in private contracts ◆ ◆. Chapter 4.

H3: public error-rate evidence remains missing

H3 is supported in the sources screened to 10 November 2026: no independent test reports a scored product's sold-category error rate with a labelled set and published denominator ◆. Failure cases, measured prevalence and disagreement between vendors answer different questions. Smith and co-authors studied publisher-side labels with partly constructed overall ratings and small overlaps; their result measures agreement, not accuracy ◆. The candidate IAS bot miss rates concern its unscored publisher tool and omit session counts ◆. Vendor-protected comparisons use the protecting classifier and cannot reveal what it missed ◆. Chapter 3.

H4: the route differs by platform

Meta, YouTube and Amazon document platform-supplied routes with different limits. TikTok's documented Open Measurement SDK integration lets a verifier's tag collect signals, contradicting H4 under its literal comparator ◆. The route statements date from 2021–2022, TikTok controls the integration, and no 2026 source opened restates the full route. Snap and LinkedIn trackers leave the firing location unstated; a named partner on other platforms supplies no missing route. Chapter 5.

H5: a narrow published safeguard exists

DoubleVerify's Supply-side line and IAS's Publisher channel show seller revenue in FY2024, with IAS's line also including Publica's ad serving ◆. The author accepts DoubleVerify's June 2024 structural-separation statement as the published safeguard that contradicts H5. It names no scope, responsible owner or enforcement for that separation and does not address the pending Nielsen-parent buyer. Ownership and seller revenue establish governance questions, not measurement bias ◆. Chapter 7.

H6: a successor depends on the process test

ISBA's March 2025 guide publishes a floor and risk tiers. The author accepts its review pledge, named overseer and update trigger as a stated change process, making H6 contradicted. Requiring a published decision procedure instead leaves it unresolved. Separately, product references to GARM do not establish which dated definitions each implements ◆. The public framework's existence therefore supplies no automatic evidence about a vendor's implemented classifier. Chapter 8.

H7: agent buying has an unresolved evidence link

The screened record contains schemas, connectors, tests and vendor production claims. Expanded reporting adds Butler/Till campaign context and quality-analysis work, but does not establish which actual verification verdict changed the buying agent's decision. No identified agreement giving that verdict financial force was found. The author retains unresolved under the stated interpretation of the rule's ambiguous named-buyer clause ◆. Schema representation alone is already too narrow a description of the work. Chapter 10.

What a buyer can use

The Verdict Quotient scores 17 products on 12 dimensions within separate cohorts. Only suitability specialists meet the band-size gate. The field values are 60 for suites, 53 for IVT specialists and 39 for suitability specialists; cross-cohort comparison is invalid. The quotient measures documented capability, without estimating accuracy. Chapter 11 exposes the cells and sensitivity tests. Chapter 12 turns the evidence into requests for scoped letters, labelled tests, data routes, correction rights and the governing contract.

INTERESTS

The author has worked at Samba TV, now branded Samba, since May 2025. Nielsen competes with it in TV measurement. DoubleVerify's agreed sale to Neptune BidCo US Inc., Nielsen's parent, remained pending at the 10 November 2026 check ◆ ◆. He co-leads AgenticAdvertising.org's Signals & Measurement working group; No Fluff Advisory operates signals and sales agents in its registry [6]. He previously worked at Verve Group, Hearts & Science and Lotame. No vendor funded, reviewed or saw the paper; every vendor was scored under the same published rules. Claude and OpenAI models performed the AI-assisted research and checks under his direction, and he checked the work. There was no human second scorer. The full interests and process statement appears in Chapter 13.

CHAPTER 1

What a verification verdict is

Ad verification sells a judgement about an ad impression against a stated condition, followed by a report or a control that acts on it. The conditions include viewability, invalid traffic, brand safety and suitability, while the delivery ranges from a list applied before bidding to measurement after an ad serves [7, 8]. This paper treats the condition, the data route and the power to act on the result as separate questions. Its comparisons follow the service a buyer purchases, with suites, invalid-traffic specialists and suitability specialists in separate cohorts.

IN SHORT

A viewable impression meets a pixel-and-time rule that establishes an opportunity to see an ad [2, 9]. Invalid-traffic filtration removes activity that falls within the Media Rating Council's general or sophisticated categories [1]. Brand safety applies an exclusion floor, while suitability lets a brand choose its tolerance within content categories [10]. A property classification and a classification of the particular image, video or audio next to an ad have different scopes [11]. Tags, server records, software development kits and platform integrations expose different evidence to the company making the judgement [12, 13, 14]. The buying-task register keeps those services apart from media sold with an embedded safety filter, publisher account protection and the platforms' own controls. A verifier's report, an exclusion, a block and a contractual right to money are separate uses, tested in Chapter 4.

The condition behind the label

A verdict needs a question narrow enough to answer. Was an impression in view for the required time? Did the measured traffic meet an invalid-traffic rule? Did the adjacent content fall into a category the advertiser excludes? The paper uses verdict for the resulting classification, measure or decision, together with the condition and scope that give it meaning. It records who measured the event, what the buyer selected and what action followed. A label without those fields leaves the purchase hard to assess.

The four metric families describe different properties of an impression. The Media Rating Council (MRC) publishes viewability and invalid-traffic requirements, while its content-level supplement addresses context, adjacency and brand safety [2, 1, 3]. A campaign can therefore need more than one test. Read a vendor's claim of quality by asking which condition it covers, which impressions were measurable and what happened to those the service could not classify. The size of the reported percentage depends on that denominator.

For display ads, MRC's desktop guidelines version 2.0 of 18 August 2015 generally require at least 50% of the pixels in view for one continuous second, with a 30% pixel exception for large display ads of at least 242,500 pixels [2]. Video generally requires at least 50% of the pixels in view for two continuous seconds while the video plays [2]. The mobile guidelines of 28 June 2016 apply to mobile web and in-app measurement and retain format-specific pixel and duration rules [9]. A buyer should record the format, environment and applicable rule before comparing rates.

MRC describes viewability as an opportunity to see, so the result alone establishes neither a person's attention nor a later response [2]. The denominator also matters within viewability. The mobile guidelines require reporting of measurable and non-measurable impressions, and distinguish a measurable impression from one that meets the viewable threshold [9]. A high rate among measurable impressions leaves a separate question about the share that could be measured. The attention paper owns attention scores and tests of attention, which this paper leaves outside its buying tasks.

Invalid traffic (IVT) is divided into general invalid traffic (GIVT) and sophisticated invalid traffic (SIVT) in MRC's June 2020 addendum [1]. GIVT includes traffic identified through routine lists or standardised checks, while SIVT calls for more advanced analysis, corroboration or human intervention [1]. The categories cover measurement exclusions as well as intentionally deceptive activity, and MRC requires services to disclose their decision rates and limitations [1]. This paper calls an impression that survives the applicable filtration classified as valid. That wording keeps the reported classification attached to the process that produced it.

Safety and suitability require a content policy as well as a measurement method. ISBA's March 2025 Responsible Media Guide distinguishes a common safety floor from brand-specific suitability and publishes high, medium and low risk tiers [10]. The floor identifies content to avoid, while the tiers express the advertiser's tolerance within categories [10]. A buyer's selected tier belongs beside the result. Otherwise two reports can differ because the advertisers asked different questions, even before considering the classifiers' performance.

What was classified

The object classified can be a domain, an app, a page or the particular content adjacent to the ad. MRC's October 2025 representation policy distinguishes property-level ad verification based on text and contextual classification from enhanced content-level work on images, video and audio [11]. Its 2018 supplement addresses discrete content measurement, metadata, adjacency, data freshness and human intervention [3]. This paper records that object explicitly, because a conclusion about a property has a different evidential reach from a conclusion about a particular video.

The content-level supplement also allows a third-party service to use validated data from the platform, subject to requirements on the data and measurement [3]. That makes the source of the information part of the result. Ask whether the verifier observed the content itself, received a content identifier, or analysed a platform-provided record. A platform integration can support a useful result, but its fields and exclusions still define what the verifier can establish. Chapter 5 records those routes by platform and metric.

The paper uses independent verification for measurement by a party other than the media seller or platform, and first-party validation for the seller's or platform's own measurement. This is a definition of the measurement relationship. It does not settle who funded a study, who pays the vendor, who owns it, or who can overturn its report. Those are separate fields. A buyer can use both types of measurement and still need to know which party controls the underlying observations.

When the verdict acts

A pre-bid service affects what the buyer considers before purchase. HUMAN's MediaGuard, for example, receives information from an OpenRTB bid request and returns an IVT prediction that the platform can use to drop the request [15]. Peer39 sells categories applied through buying platforms, and Proximic describes pre-bid inventory filters for media buyers [16, 17]. The decision is made against information available at that point. A later report about the served impression answers a different question.

Post-bid measurement records an impression after the purchase or serve event. HUMAN's FraudSensor collects device, network and behavioural signals through detection tags and returns IVT classifications in reports [14]. Pixalate documents tags and server-to-server event submission for its post-bid Analytics product, alongside separate pre-bid data feeds [12]. Post-bid blocking is another action again. IAS describes pre-bid filtering and post-bid blocking for content outside a brand's settings, while Adloox describes page-context blocking and exclusion [8, 18].

The timing should stay attached to the metric. A product can sell an IVT prediction before bidding, viewability measurement after delivery and suitability lists for a platform to execute. A buyer should ask which component produced each report and whether blocking was enabled for that placement. Merely finding the vendor's tag in an ad does not settle those settings. The useful record pairs the delivered metric with its collection route, stage and permitted action.

How the evidence reaches the verifier

A tag runs code or sends a request when the ad or page loads. HUMAN documents detection tags that gather device, network and behavioural signals, while Pixalate documents JavaScript tags, pixels, video verification scripts and server-submitted events [14, 12]. Each route has a distinct integration requirement. In the buying record, keep the tag's event separate from any model or classification subsequently applied to it. That makes it possible to ask which observations a missing or blocked tag would remove.

The Open Measurement Software Development Kit (OM SDK), maintained by IAB Tech Lab, provides measurement signals across supported apps and video environments [13]. The Tech Lab's May 2024 CTV release describes signals such as device type and TV-off detection and names platform-specific integrations [13]. Support for the SDK therefore needs an implementation record for the relevant player or device. Chapter 6 examines the CTV and in-app routes, including the distinction between an ad server's event and a device observation.

A server record can carry an impression identifier, time, app or device information supplied by another system. Pixalate explicitly supports server-to-server submission for platforms that cannot run a client-side tag [12]. A crawler and classifier instead inspect content and assign categories. IAS describes analysis of page content and app metadata, supplemented by human review, and Channel Factory describes video classification used to build YouTube inclusion lists [8, 19]. Record the input and the output separately so that a source's description of a collection method does not become evidence of classification accuracy.

Platform feeds and application programming interfaces (APIs) form another route. Sightly says its YouTube suitability classifications use contextual signals from the YouTube Data API and that Google does not review those classifications [20]. Its self-service list builder is consequently a different service from a platform's own safety control or a report measuring specific ad impressions [20]. The table below preserves that distinction between selecting content and measuring delivery.

What the assurance covers

The paper reserves accreditation for MRC's scoped service record. MRC's representation policy distinguishes property-level from content-level accreditation and requires vendors to state the relevant capabilities and limits [11]. TAG's Brand Safety Certified programme instead certifies a company's material operations in a stated geography, at entity level [21]. The issuer, programme, entity and scope should therefore accompany the seal. Treat a vendor's own statement about compliance as self-attestation until the named body's record establishes what assurance it granted.

These assurance records answer a further question from the metric itself. A buyer still needs to know which product, setting and environment the purchase uses. A company-level seal cannot fill the product's missing method, and a method document cannot supply a test's error rate. The scoring rules give those attributes separate homes so that the same assurance does not earn credit throughout a row.

The service being bought

The register groups products by the output purchased. Cohort A contains suites sold to buyers for viewability, GIVT and SIVT, and at least one safety or suitability service across two or more environments. Cohort B contains impression-level IVT measurement or filtering for advertisers, agencies or trading platforms. Cohort C contains content safety or suitability classification sold for verification, exclusion or blocking. A product has one row for its buying task, even where it bundles several delivery components.

The test for a suitability product is whether the buyer can purchase the classification or control apart from managed media. Pixability's Self offer lets advertisers activate contextual and suitability data through their own Google Ads or Display & Video 360 account [22]. Sightly's self-service offer includes YouTube inclusion and exclusion lists, and Channel Factory's August 2026 terms describe software accounts and paid plans alongside its services [20, 23]. These are recorded separately from the companies' managed campaigns and curated media offers.

HUMAN and Pixalate remain in the IVT cohort under this paper's primary-output rule, even though their menus also describe viewability or other quality signals [15, 24]. Peer39's suitability row and its Adloox verification suite remain separate purchased services, with Peer39's June 2026 acquisition announcement describing the addition of viewability and fraud detection [25]. The method applies to the named service. Ownership alone supplies no missing capability, and one row's evidence cannot fill another row's cells.

Q2 records the products in the scoring population below. Product capabilities, delivery and environment descriptions are vendor documentation, unless the evidence column identifies a filing. The table describes the offer and leaves accreditation coverage to Chapter 2. Owner unknown means ultimate control was not established in the register, rather than an inference of independence. DSP means demand-side platform, the software through which a buyer bids. CTV means connected television. MFA means made-for-advertising. Environment labels in this table describe the vendor's offer, with each metric still requiring its own scope check.

PRODUCT	RECORDED OWNER OR OPERATOR	COHORT	METRICS SOLD	METHOD AND STAGE	BLOCKS OR REPORTS; VENDOR-DESCRIBED ENVIRONMENTS	EVIDENCE
DoubleVerify advertiser verification	DoubleVerify Holdings, public company at the baseline read	A	IVT, viewability, safety and suitability	Pre-bid segments and post-bid measurement	Avoidance, blocking and reporting; web, apps, CTV and platform placements	[7, 26]
IAS advertiser verification	Novacap affiliates, through Igloo Group Parent	A	IVT, viewability, safety and suitability	Tags, classification and platform integrations, pre-bid and post-bid	Filtering, blocking and reporting; web, apps, CTV, audio and platforms	[8, 27, 28]
Adloox verification	Peer39	A	IVT, viewability, safety and suitability	Measurement, page analysis and DSP segments, pre-bid and post-bid	Exclusion, blocking and reporting; web, mobile and video/CTV offer	[18, 29, 25]
Protected by Mediaocean	Mediaocean	A	IVT, viewability and suitability	Detection and verification suite, post-bid with pre-bid controls	Filtration, controls and reporting; desktop, mobile web, apps and CTV, metric-specific	[30, 31]
mFilterIt Ad Traffic Validation	mFilterIt; control unconfirmed	A	IVT, viewability, MFA, safety and suitability	Pre-bid avoidance and post-bid validation	Avoidance and reporting; display, video/CTV and apps	[32]
HUMAN Ad Integrity Suite	HUMAN Security; control unconfirmed	B	GIVT and SIVT, with other quality signals	Bid-request predictions and post-bid tags	Platform-applied filtering and reporting; environment-specific implementations	[15, 14]
Pixalate ad fraud protection and analytics	Pixalate; control unconfirmed	B	IVT and display viewability, with quality analytics	Pre-bid feeds and post-bid tags or server events	Blocking feeds and reporting; websites, apps and CTV	[24, 12]
Fraudlogix IVT detection	Fraudlogix; control unconfirmed	B	GIVT and SIVT, with quality flags	Post-bid pixel analytics and pre-bid IP lists	Platform-applied blocking and reporting; high-volume real-time bidding traffic	[33]
Escalated IVT scoring	Escalated Solutions; control unconfirmed	B	IVT and viewability	Pre-bid scoring and post-bid monitoring	Blocking and reporting; display, video, apps and CTV	[34, 35]
FouAnalytics	Founder-operated; control unconfirmed	B	In-ad IVT analytics	JavaScript tag and post-bid classification	Reporting; ads and websites where its tag executes, separate landing-page measurement for restricted platforms	[36]

PRODUCT	RECORDED OWNER OR OPERATOR	COHORT	METRICS SOLD	METHOD AND STAGE	BLOCKS OR REPORTS; VENDOR-DESCRIBED ENVIRONMENTS	EVIDENCE
Zefr brand safety and suitability	Zefr; control unconfirmed	C	Content safety and suitability	Classification, pre-impression exclusions and post-campaign measurement	Exclusion and reporting; social-platform content	[37]
Peer39 brand safety and suitability	P39 Tech; ownership self-description in register	C	Safety and suitability categories	Pre-bid content classification and post-buy reporting	Buyer-applied exclusions and reporting; web and CTV, by product scope	[16, 38]
Proximic by Comscore	Comscore	C	Safety, suitability and misinformation filters	Pre-bid contextual segments	Buyer-applied exclusion; desktop, mobile, CTV and podcasts	[17]
illumaSentinel	Illuma Technology; control unconfirmed	C	Page-level safety and suitability	Pre-bid classifier, activated in The Trade Desk	Avoidance and reporting; open-web pages	[39]
ViewIQ	Channel Factory; Truelink investment, control unconfirmed	C	YouTube safety and suitability	Pre-bid classification and inclusion/exclusion lists	Buyer-platform execution and reporting; YouTube	[19, 23, 40]
Pixability brand suitability	Pixability; control unconfirmed	C	YouTube suitability	Contextual segments and inclusion/exclusion lists before bidding	Buyer-account activation and reporting; YouTube	[22]
Sightly Brand Mentality	Sightly; control unconfirmed	C	Suitability SmartLists	YouTube API-based classification and list building	Inclusion/exclusion targeting; YouTube and programmatic offers, with separate CTV deals	[20]

Mantis Sentinel, a monitoring tag from Reach's Mantis, is listed rather than scored: its advertiser page describes the tag and its page-level output but gives neither separate purchase terms nor a live or beta status, which the scored-set rule requires.

A product comparison should hold the purchased task constant. Start with the output the buyer receives, then check the metric, environment, timing and action that make it useful. A publisher's account-protection tool and an advertiser's measurement service can both detect suspicious activity while answering different buying needs. The row boundary prevents that shared vocabulary from turning into a comparison the evidence cannot sustain. It also prevents a broad suite from appearing again as several specialist competitors.

For the same reason, the paper separates content inclusion from an observation of delivery. A list can express a brand's intended policy, while a report can describe what appeared during a campaign. To connect the two, ask for the settings applied, the period they covered and the relevant delivery record. That is a request for evidence about the named use. It should travel with the buyer's question before a coverage claim or summary score is used in procurement.

Who pays, and what the totals contain

The paying customer can be an advertiser, agency or trading platform, while publishers also buy measurement and quality tools. DoubleVerify's FY2025 filing separates advertiser

Activation and Measurement from Supply-side revenue, and IAS's FY2024 filing separates Optimization, Measurement and Publisher revenue [26, 27]. Those labels describe different commercial perimeters. The buyer of a report and the party whose inventory it assesses can therefore sit in different contractual relationships with the same vendor.

DoubleVerify filed worldwide revenue of US$748.3 million for calendar 2025, including US$71.3 million from Supply-side customers [26]. IAS filed worldwide revenue of US$530.1 million for calendar 2024, including US$76.5 million in its Publisher channel ◆. The IAS line includes Publica's ad serving, while DoubleVerify's Supply-side line is a customer grouping ◆. These company totals have stated years and scopes. They cannot establish the size of a common verification market or a market share without a comparable denominator.

Revenue also uses different units from product activity. DoubleVerify charges advertisers a measured transaction fee tied to the ads measured, while IAS describes recognising revenue from a fee per thousand impressions multiplied by measured volume [26, 27]. A buyer should keep the vendor's billing base beside its detection rate and its reported volume. Chapter 7 examines customer types, owners and safeguards, with the transaction record dated to the evidence read.

The register's boundaries

The product register holds 45 records, of which 19 are listed and 9 historical. A listed record documents a boundary decision, rather than a score of zero. Anura's Ad Tag Protect, for example, protects a publisher's own advertising account by withholding ads from flagged traffic, and its visitor-scoring offer addresses site visits [41, 42]. The scored IVT task instead requires a documented impression product serving a buyer or trading platform.

Contextual media and curated deals have their own papers in this series. GumGum's contextual advertising page and Seedtag's safety policy describe protections within their media offers [43, 44]. This paper requires a separately bought verification, exclusion or blocking service before treating that signal as a suitability row. Attention scores, content-rights protection and a platform's own advertising controls remain separate buying tasks. A broad use of the word safety cannot decide the product's cohort.

Historical records preserve earlier names without creating additional live competitors. Oracle's August 2024 end-of-life notice names Moat Analytics and Oracle Contextual Intelligence and ends all Oracle Advertising services on 30 September 2024 [45]. DoubleVerify's filing records its purchases of Meetrics and OpenSlate [26]. The register also keeps Method Media Intelligence as historical, with current commercial status unresolved after the domain served an unrelated IT-services page when screened [46]. A dead or redirected page is an access and status observation. It cannot, alone, establish a corporate dissolution.

The questions a verdict has to survive

The author's essay Where the Surplus Settles argues that advertising verification has little pricing power when its verdict does not change a price [4]. The Risk You Can Price identifies

a missing adjudicator for disputed conversion counts, a narrower claim about cross-company outcome settlement [5]. Here those are rival framings to test. They set questions about contractual authority, payment and redress, rather than conclusions about every verification product. Chapter 4 examines the actual terms that could support or defeat them.

H1a asks how much of the sold product-by-metric-by-environment grid holds MRC accreditation. H1b asks whether public accreditation claims stay inside the applicable record. Chapter 2 puts the service, metric, environment, format, placement, language and named standard beside each scope. That establishes what the assurance covers before using a company logo as shorthand.

H3 asks what independent tests establish about detection. Chapter 3 keeps a labelled error-rate test apart from a failure example, a prevalence study or agreement between classifiers. H4 asks how third parties obtain platform data and what they can do with it. Chapter 5 records the route for each partner, while Chapter 6 examines the device and server dependencies in CTV and apps.

H5 asks who pays the verifiers and what safeguards they publish, in Chapter 7. H6 asks whether a maintained common suitability framework exists after GARM, in Chapter 8. Chapter 9 keeps allegations, letters, procedural decisions and findings on the merits in their own records. They inform the questions a buyer asks and remain outside the Verdict Quotient.

H7 asks whether agent buying makes a verification result binding in an executed transaction, in Chapter 10. The Verdict Quotient compares documented capability within each buying task in Chapter 11, and Chapter 12 turns the evidence into requests for methods, scope, authority and correction rights. Chapter 13 records the evidence dates, verifier passes and unresolved items. Status is as read in November 2026; a recheck on 11 December 2026 is published as dated corrections.

CHAPTER 2

What MRC accreditation certifies

The Media Rating Council (MRC) accredits measurement processes against named standards and scopes. Its letters publish no product error rate ◆ ◆. On the lists read on 10 November 2026, content-level brand safety and suitability accreditation is confined to YouTube ◆. On this paper's reading, Peer 39's wording implies pre-bid accreditation beyond the listed scope, contradicting H1b ◆. Every status is dated to its read and will be rechecked on 11 December 2026.

IN SHORT

MRC's letters certify conformance with named standards and publish no error rate ◆ ◆. Its standards and guidance read on 10 November 2026 span February 2012 to July 2026 ◆. TAG's four July 2026 certification programmes cover operating companies ◆. Only Google's first-party controls and Zefr's third-party reporting hold content-level accreditation, both on YouTube with different placement and language limits ◆. GumGum withdrew and its accreditation was revoked in September 2025 ◆. IAS's property-level row has no current letter basis ◆. H1b is contradicted on this paper's interpretation of Peer 39's wording and a literal reading of Protected by Mediaocean's display-impression claim on CTV ◆ ◆. H1a is contradicted on the merged documented grid, with the denominator limits stated below ◆.

What an accreditation certifies

MRC accredits one service's measurement process at a time. It is an industry-funded body whose members include advertisers, agencies, media sellers and measurement firms, and its audits are performed by CPA firms that MRC engages and the audited service pays for [47]; EY appears on the copy lists of its letters to HUMAN (formerly White Ops), Pixalate and Zefr ◆ ◆ ◆. Each letter certifies that the service's reported statistics comply with the standards it names and with the procedures the service described. The audit tests conformance and reports no detection or error rate ◆ ◆. Zefr's release of 9 March 2026 glosses its content-level accreditation as confirming "validity, reliability, and governance"; MRC's letter uses neither of the first two words ◆.

MRC's response to Adalytics on 11 April 2025 requires impression-level GIVT filtration after serving, while pre-bid blocking is optional ◆. Pre-bid audits test production and accuracy of supplied files and databases, while the platform decides whether to bid ◆. HUMAN's accreditation covers pre-bid detection rather than prevention of serving, and MRC says it has no billing insight ◆ ◆.

MRC's February 2024 policy makes the letter the authoritative scope reference ◆. The digital list read on 10 November 2026 held 41 services without an update stamp ◆. Superseded letters may disappear, and quarterly reports can omit continuations ◆. Withdrawal records an exit from the audit process without establishing an accuracy failure ◆.

Accreditation sits beside two neighbours that Chapter 1 defines: TAG certification, which covers an operating company rather than a product, and vendor self-attestation, such as a "Brand Safety" label from a vendor MRC does not audit ◆ ◆.

The standards and their versions

The table lists issuer documents read on 9 and 10 November 2026 and the letters or seals that cite them ◆ ◆. Guidance dates, accreditation dates and software releases describe different events ◆. The July 2026 AI guidance was only scanned in the earlier pass, so this chapter draws no compliance conclusion from it ◆.

ISSUER	DOCUMENT	VERSION AND DATE	WHAT IT COVERS	NAMED IN THE LETTERS OR SEALS OF	EVIDENCE
MRC	Viewable Ad Impression Measurement Guidelines (desktop)	Version 2.0, 18 August 2015	Viewable-impression thresholds	Every viewability letter	◆
MRC	Mobile Viewable Ad Impression Measurement Guidelines	Final, 28 June 2016	Viewability on mobile web and in-app	Every viewability letter	◆
MRC	Invalid Traffic Detection and Filtration Standards Addendum (IVT 2.0)	Issued October 2015; updated June 2020	Detection and filtration of GIVT and of sophisticated invalid traffic (SIVT)	DoubleVerify, IAS, HUMAN, Pixalate, Protected by Mediaocean and Adloox; Zefr's YouTube viewability letter for the GIVT provisions only	◆ ◆ [1]
MRC	2024 IVT Interim Updates	Final, 24 April 2024	Privacy; spoofed CTV domains and bundle IDs; property-level considerations for SIVT vendors	Cited as interim guidance in most 2025 and 2026 letters	◆ [48]
MRC	Statement on IVT requirements and processes for pre-bid	11 April 2025	GIVT filtered after serving; pre-bid filtration not required; what pre-bid audits cover	A statement, not a standard	◆
IAB and MRC	Guidelines for the Conduct of Ad Verification	14 February 2012	Property-level ad verification	DoubleVerify; Protected by Mediaocean's property-level letter; Zefr's and Google's content-level letters	◆ ◆
MRC, with the American Association of Advertising Agencies (4A's), the Association of National Advertisers (ANA) and IAB	Supplement: Enhanced Content Level Context and Brand Safety	Final, September 2018	Content-level brand safety; reporting granular enough for buyers' own suitability settings; platform data access	Protected by Mediaocean's property-level letter; Zefr's and Google's content-level letters	◆ ◆
MRC	Server-Side Ad Insertion and OTT Guidance	Final, August 2021	Measurement of server-inserted ads and of video delivered to TVs	Protected by Mediaocean's property-level letter	◆ [49, 50]

ISSUER	DOCUMENT	VERSION AND DATE	WHAT IT COVERS	NAMED IN THE LETTERS OR SEALS OF	EVIDENCE
MRC	Policy for Accreditation Representations by MRC-Accredited Services	15 February 2024	How accredited services may describe their accreditation	A policy, not a standard	◆
MRC	Policy for Property-Level Ad Verification Representations	18 October 2025, with a six-month grace period	"Brand Safety" labels and disclosures for property-level services	A policy, not a standard	◆
MRC	Guidance on Existing Standards for AI Use in Media Measurement	8 July 2026	Maps AI-related requirements in seven MRC documents; new AI standards in phases to early 2027	Guidance	◆
IAB Tech Lab	Open Measurement SDK, JavaScript clients	1.6.12, 2 November 2026	Measurement interface for in-app, web video and CTV; device attestation since 1.6.0 of 29 October 2025	DoubleVerify's letter (in-app viewability including OM SDK traffic); required by TAG's Certified Against Fraud 11.0 for verifiers under its device-attestation guidance	◆ ◆
IAB Tech Lab	ads.cert 2.0	Page last updated 14 February 2024	Authenticated Connections published; Authenticated Delivery and Authenticated Devices marked future	None found	◆ ◆
TAG	Brand Safety Certified Guidelines	5.0, July 2026	Brand-safety certification of the operating company	Seals held by DoubleVerify, IAS, Adloox, Channel Factory and illuma	◆ ◆
TAG	Certified Against Fraud Guidelines	11.0, July 2026	Anti-fraud certification of the operating company	Seals held by DoubleVerify, IAS, HUMAN, Pixalate, Protected by Mediaocean and Adloox	◆ ◆
TAG	Certified Against Malvertising Guidelines	7.0, July 2026, renamed from Certified Against Malware	Anti-malvertising certification of the operating company	No seal found among the companies searched	◆ [51]
TAG	Certified for Transparency Guidelines	Version 2, July 2026	Transparency certification of the operating company	Seals held by DoubleVerify, IAS and Protected by Mediaocean	◆ ◆

MRC's standards page read on 10 November 2026 names no separate suitability guideline [52]. The 2018 supplement requires discrete content reporting for buyers' risk settings, points to the 4 A's framework, and requires direct measurement or validated platform data for third-party content-level accreditation ◆. It also requires disclosures of expected errors, sampling uncertainty and technical limits, while making false-positive likelihood optional [3]. These disclosures do not supply a public product error rate [3]. Chapter 5 covers data routes and Chapter 8 the frameworks.

Two MRC policies govern what accredited services may say. The February 2024 policy asks them to state what is and is not accredited, bars claims for anything not accredited, including services under review, and restricts "first" and "only" claims ◆. The policy of 18 October 2025, addressed to the services MRC audits, bars accredited property-level services from positioning themselves or their metrics as "Brand Safety", or as content-level, without content-level capability in the environment, after a six-month grace period that by the paper's computation ran to about 18 April 2026; suitability language stays allowed with

disclosures ◆. It requires property-level vendors to disclose that images, video and audio are excluded, among other limits, and says MRC will not consider new property-level submissions without commitments to content-level capability [11]. It does not bind vendors MRC does not audit, and MRC does not say whether it reaches an audited vendor's unaudited products, such as the "Brand Safety" products Zefr sells off YouTube, so that question is unresolved ◆ ◆.

In the MRC records screened to 10 November 2026, no decision applying the October 2025 representation policy was found ◆. DoubleVerify's letter of 13 April 2026 still uses Brand Safety metric names, while its one-sheet uses Brand Risk Floor Avoidance ◆. Protected Media's pre-policy letter uses Brand Safety names with a property-level caveat [50]. These discrepancies support no compliance finding ◆.

The Tech Lab's ads.cert 2.0, a protocol for authenticating programmatic counterparties, has a reference repository whose last commit is dated 8 March 2023, and the pages read publish no adoption count for it or for OM SDK ◆ ◆.

TAG certifies companies

TAG reissued all four of its certification guidelines in July 2026: Brand Safety Certified 5.0, whose change log runs back to version 1.0 of September 2020, when the Digital Trading Standards Group of the UK's former Joint Industry Committee for Web Standards (JICWEBS) merged into TAG's programme; Certified Against Fraud 11.0, which makes intermediaries use PaymentID and drops a billing-notification requirement for server-side ad insertion (SSAI) on CTV; Certified Against Malvertising 7.0, renamed from Certified Against Malware; and Certified for Transparency version 2 ◆. TAG certifies the operating entity, not a product, and its brand-safety guidelines require a content verification service to disclose its capabilities and to review complaints of wrongful classification objectively ◆ [21].

TAG's registry read on 10 November 2026 lists DoubleVerify and IAS with global fraud and brand-safety seals confirmed through independent validation and TAG-validated transparency seals for North America and Europe ◆. Both appear on TAG's Platinum page ◆. Zefr and Escalated were enrolled without seals, while searches for Peer 39, Comscore and Fraudlogix returned no entry under the names tried ◆. TAG seals do not enter H1a's MRC numerator.

MRC's decisions, 2023 to 2026

Between January 2023 and September 2026, MRC's quarterly reports record grants that widened viewability, invalid-traffic and platform reporting, and revocations of verification services or platforms' brand-safety controls in March 2023, November 2023, June 2024 and September 2025, each after a withdrawal from the process or a product's deprecation [53, 54, 55, 56]. The table lists the grants, revocations, denials and removals, and the continuations on which current scopes rest.

DECISION	SERVICE	ACTION	SCOPE	EVIDENCE
January 2023	GumGum (Verity API)	Grant, continue, deny	Content-level classification and URL reporting granted on CTV (English); desktop and mobile web continued; in-app denied	[53]
March 2023	Facebook Instant Articles (Meta, first-party)	Revoke	Content-level brand-safety controls on Instant Articles, a deprecated product; in-stream video remained	[53]
May 2023	IAS	Grant	CTV video tracked ads and rendered and viewable impressions, certified traffic only	[57]
September 2023	Adloox	Grant	In-app display and video viewability, exposure metrics and decision rate	[58]
November 2023	Method MI	Revoke	Rendered display ad impressions, after a hiatus, on withdrawal from the audit	◆
March 2024	DoubleVerify	Grant, continue, deny	Granted: Serbo-Croatian (domain), Korean and Hindi (CTV), CTV video viewability including SIVT, pre-bid decisioning; denied: pre-bid Authentic Attention	[59]
March 2024	IAS (YouTube through Google Ads Data Hub)	Grant	Third-party processing of YouTube video impressions and viewability on desktop, mobile web and in-app	[59]
April 2024	IAS	Grant	SIVT filtration and reporting on CTV	[55]
June 2024	Oracle Moat Analytics, Moat's YouTube and Meta reporting, Oracle Contextual Intelligence	Revoke	On withdrawal from the process and planned deprecation	[55]
Letter of 9 October 2024	IAS	Continue, in no quarterly report	Property-level ad verification for display and video, English only, on desktop and mobile web, with the platform's other metrics	◆ ◆
April 2025	Pixalate	Continue	Ad counts, impressions and SSAI metrics with SIVT on desktop, mobile web, in-app and CTV; display viewability on the first three	◆
May 2025	GumGum	Continue	Content-level brand safety for text, image and video, English, on desktop, mobile web and CTV, open web	[60]
June 2025	Meta (first-party)	Grant, continue	Facebook and Instagram Feed controls and reporting granted (English, Spanish); Facebook in-stream continued (37 languages); misinformation excluded	[60]
September 2025	Protected Media	Grant	Property-level ad verification on desktop and mobile web, English, Spanish and Portuguese	◆ [56]
September 2025	GumGum	Revoke	Content-level brand safety and URL reporting on the open web, on withdrawal from the process	◆ [56]
September 2025	Meta (first-party)	Revoke	Feed and in-stream content-level controls and reporting, on withdrawal from the process	◆ [56]
November 2025	IAS reporting for Amazon's demand-side platform (DSP), server to server	Grant	Impressions, viewability and tracked ads with SIVT on desktop, mobile web and in-app; Fire TV and Twitch excluded	[61]
November 2025	Zefr (YouTube through Google Ads Data Hub)	Grant	Video impressions and viewability on desktop, mobile web and in-app; skippable, non-skippable and bumper ads	◆ [61]
December 2025	TikTok	Removed from in-process	Lack of intent to engage in a third-party audit	[61]
January 2026	IAS Attention Measurement; IAS In-Game Measurement	Removed from in-process	Lack of intent to engage in a third-party audit	◆
January 2026	Zefr (YouTube content-level)	Grant	English in-stream video and music video on desktop, mobile web, in-app and CTV, with exclusions	◆ [62]
April 2026	DoubleVerify	Grant, continue	TikTok in-app video viewability with GIVT and SIVT; French on CTV and five domain-level languages; attention reporting (the attention paper's subject)	[63] ◆

DECISION	SERVICE	ACTION	SCOPE	EVIDENCE
April 2026	HUMAN	Grant, continue	Viewability on desktop, mobile web and in-app granted; pre-bid and post-bid GIVT and SIVT on four environments continued	◆
April 2026	Google (YouTube, first-party)	Grant, continue	Shorts granted; in-stream continued	[63]
June 2026	IAS	Grant, continue	Pre-bid viewability and IVT-avoidance segments on desktop, mobile web and in-app granted; platform continued "excluding Ad verification"	◆ ◆
June 2026	Peer39 Adloox (formerly Scope3 France Adloox)	Continue	Display and video tracked ads, impressions and viewability on desktop, mobile web and in-app; SIVT on desktop and mobile web	◆
September 2026	Protected Media	Continue	Same scope as April 2026	[64]

Two gaps shape the record. IAS's continuation of 9 October 2024, the last letter found granting its property-level ad verification, appears in no quarterly report and is read from an Internet Archive copy, because the live file returns HTTP 404 ◆ ◆. Pixalate's scope still rests on its letter of 8 April 2025: none of the five reports from July 2025 to September 2026 records a Pixalate decision, and a missing entry is not evidence that no continuation occurred ◆.

Accreditation by product

The scored products hold most of their MRC accreditation in viewability and invalid traffic, as the figure shows. Nine of the scored products in the figure have no MRC entry at all: mFilterIt among the suites; Fraudlogix and FouAnalytics among the invalid-traffic specialists; and Peer39's own product, Proximic by Comscore, illumaSentinel, ViewIQ, Pixability and Sightly's Brand Mentality among the suitability specialists [47, 65]. Escalated has nothing accredited, but MRC lists its audit as in process as of June 2026 ◆.

MRC STATUS BY SCORED PRODUCT AND METRIC FAMILY, 10 NOVEMBER 2026



D desktop, MW mobile web, App in-app, CTV connected TV; all 4 = all four; +n platf. = the number of large platforms where a platform-specific service is also accredited (YouTube via Ads Data Hub, Meta, TikTok, Amazon server to server). IAS's property-level row is listed on MRC's page without a current letter (the last letter granting it is dated 9 October 2024). Accreditation certifies a measurement process for the listed metric, environment and scope; its audits can include accuracy testing, but the letters publish no product error-rate estimate. Read on MRC's accredited and under-review lists and letters, 10 November 2026; status-sensitive and rechecked at the freeze.

> **Nine of seventeen scored products have no MRC entry, and only Zefr holds content-level accreditation.** MRC status for each scored product, by metric family and environment, from MRC's accredited and under-review digital lists and the letters linked from them, read on 10 November 2026. A product with no MRC entry may still sell the family; the matrix scores accreditation against what each product sells (chapter 11). Content-level accreditation outside YouTube is held by no scored product.
>
> Units: status category · Scope: the 17 scored products of the product register, MRC digital accreditation, US · Sources [101] [47] [82] [65] [63] [67] [147] [55] [66] [148] [58] [149] [64] [50] [56] [150] [151] [60] [152] [62] [153] [61]
>
> Status-sensitive: MRC's page carries no update date. "No MRC entry" is bounded to MRC's digital accredited and under-review lists screened that day.

Content-level brand safety and suitability accreditation is held by two services, both for YouTube, with different placements ◆. Google's is first-party validation of its own controls, not third-party verification: in-stream video and, since an April 2026 decision, Shorts, on desktop, mobile web and in-app according to the list ◆. Zefr's is third-party labelling and reporting of English-language in-stream video and music video on desktop, mobile web, in-

app and CTV, with Shorts excluded ◆. On the list read on 10 November 2026, no open-web service holds content-level accreditation, and Zefr is the only scored product that does; both content-level letters cite the 2018 supplement ◆ ◆ ◆. Zefr also sells content-level classification on Meta, TikTok, Snap and Reddit, none of it accredited ◆. IAS's YouTube brand-safety audit is listed as in process as of April 2023, Meta's pipeline that supplies brand-suitability data to third parties is under review, and MRC's under-review list names no DoubleVerify service ◆ ◆ [65].

In May 2025 MRC continued GumGum's English-language content-level accreditation for open-web text, image and video on desktop, mobile web and CTV [60]. GumGum's content-level accreditation was revoked in September 2025 after the contextual paper recorded it; this paper corrects that record rather than repeating it. The revocation followed GumGum's withdrawal from the accreditation process, and Meta's first-party content-level accreditation was revoked the same month on the same ground ◆ [56]. A withdrawal is not a demonstrated accuracy failure ◆.

DoubleVerify's and Protected by Mediaocean's accreditations in this area are property-level ad verification, which MRC's 2025 policy treats as classification of properties from text or keywords, with images, video and audio excluded [11]. DoubleVerify holds it on desktop, mobile web, in-app and CTV under its letter of 13 April 2026, in about 176 languages for binary classification, 56 at domain level, 41 at page level and ten on CTV, and its own one-sheet states that this is not content-level brand safety ◆ ◆. Protected by Mediaocean holds it on desktop and mobile web only, from an initial grant of 25 September 2025 covering display reporting and blocking and video reporting in English, Spanish and Portuguese ◆.

IAS's property-level row has no letter on MRC's site. The list shows IAS with property-level ad verification on desktop and mobile web, but the last MRC letter whose body grants it is the continuation of 9 October 2024 (display and video, English only), which the list linked in every capture read from 28 April 2025 to 16 June 2026 ◆. The letter of 9 July 2026 that the list now links omits ad verification, MRC's report for April to June 2026 says that continuation excluded it, and MRC's report for July to September 2026 records no IAS action ◆ ◆. Whether the accreditation lapsed, continues on an unreported cycle or is shown in error is unresolved, and for H1a the row has no current letter basis ◆.

Pre-bid detection of invalid traffic is accredited for three services: DoubleVerify and HUMAN on desktop, mobile web, in-app and CTV, and IAS on desktop, mobile web and in-app from an initial grant in June 2026, although MRC had said on 11 April 2025 that IAS was then neither audited nor accredited for it ◆. Adloox, owned by Peer39 since June 2026, holds SIVT filtration on desktop and mobile web only ◆. Peer39's only accredited service is that Adloox service, scored as its own row because evidence is not inherited between rows, and Comscore's own page claims MRC accreditation only for Comscore TV ◆. HUMAN's June 2026 suite sells brand safety and suitability with no MRC accreditation or listed audit for it ◆, and Pixalate's brand-safety ratings of apps and sites are property-level and neither accredited nor under review ◆.

On CTV, SIVT filtration is accredited for eight services on MRC's list, five of them scored products, and video viewability for two, DoubleVerify and IAS, both limited to certified apps; Chapter 6 reads that record ◆ ◆ ◆. On the nine large platforms, third-party measurement of a platform's own placements is accredited on four and third-party content-

level classification on one, YouTube, through Zefr; Chapter 5 reads what those accreditations rest on ◆.

Every accreditation above is narrowed by its letter. The table gives the placement, language, data and environment limits that H1a's units inherit.

PRODUCT	SERVICE AND LETTER	WHAT THE LETTER NARROWS	EVIDENCE
DoubleVerify	Quality Analytics Platform, 13 April 2026	On CTV, only certified apps count toward video viewable impressions and Fully On-Screen metrics, while other CTV viewability metrics may include tested, non-certified apps; no display viewability on CTV	◆
DoubleVerify	Same letter	Property-level languages: about 176 binary, 56 at domain level, 41 at page level, 10 on CTV	◆
DoubleVerify	Same letter	Pre-bid brand safety and suitability segments on desktop, mobile web and in-app only; CTV pre-bid limited to fraud and IVT, Fully On-Screen, app-store categories and site and app lists	◆
DoubleVerify	Meta, YouTube and TikTok reporting	Meta: Facebook and Instagram feed placements, net of GIVT; YouTube: in-stream and bumper through Ads Data Hub, net of GIVT, with Shorts, in-feed and masthead out of scope; TikTok: in-app video with GIVT and SIVT filtration	◆
IAS	Campaign Monitor, Firewall and Viewability, 9 July 2026	CTV video metrics and SIVT for certified applications only; metrics under the mobile rich-media ad interface (MRAID) out of scope	◆
IAS	Same letter	Pre-bid viewability and IVT-avoidance segments (IP and user-agent segments, Fraud IP List) on desktop, mobile web and in-app; no CTV	◆
IAS	Property-level row on the list	No current letter; the last grant, of 9 October 2024, covered display and video, English only, on desktop and mobile web	◆
IAS	Amazon DSP server-to-server reporting, 7 November 2025	Amazon Retail, IMDb, Amazon Publisher Services and third-party exchange supply; Fire TV and Twitch excluded	[66]
IAS	Meta reporting, 9 July 2026	Placements audited for Meta's own measurement, many excluded; net of GIVT; IAS cannot identify IVT inside Meta independently	[67]
Adloox	Peer39 Adloox, 29 June 2026	SIVT filtration on desktop and mobile web only; in-app metrics net of GIVT; no pre-bid metric	◆
Protected by Mediaocean	Protected Media, 22 September 2026	On CTV, video impressions only, by non-SDK measurement; no CTV viewability	◆
Protected by Mediaocean	Property Level Ad Verification, 25 September 2025	Desktop and mobile web; display reporting and blocking, video reporting only; English, Spanish and Portuguese; gambling and misinformation categories, IVT ad verification, custom bundle lists and rules, sentiment and non-matching language excluded	◆
HUMAN	Ad Measurement Platform, 14 April 2026	Viewability on desktop, mobile web and in-app only; the list's SIVT row applies to backend detection and IVT predictions; pre-bid accreditation covers detection, not prevention	◆ ◆ ◆
Pixalate	Analytics platform, 8 April 2025	Display viewability on desktop, mobile web and in-app only; video viewability and pre-bid SIVT in audit as of September 2025; no CTV viewability accredited or under review	◆ ◆
Zefr	YouTube content-level, 20 January 2026	English-language in-stream video and music video; Shorts, YouTube TV, Google TV, Masthead, Google Video Partner videos and volumetric data excluded	◆
Zefr	YouTube through Ads Data Hub, 6 November 2025	In-stream and bumper video on desktop, mobile web and in-app; GIVT filtration only; in-feed, masthead, Shorts and non-YouTube Demand Gen inventory excluded	◆
Google (first-party, not scored)	YouTube Brand Safety, 28 May 2026	Google Ads and Display & Video 360 only; in-stream and Shorts; video discovery, masthead, YouTube Kids, non-YouTube inventory and live streams excluded; no CTV on the list	◆ [68]

Whether vendors' claims stay within scope

H1b is contradicted on this paper's reading of what the pages imply. Its rule requires at least one current claim that states or implies an unlisted metric, environment, format or placement. The Peer 39 and Protected readings below meet that test only under the interpretations stated beside them ◆ ◆.

Peer 39 places accredited viewability measurement inside its description of the pre-bid signal ◆. The paper reads this as implying pre-bid accreditation, which MRC's separate post-serve Adloox letter does not grant ◆. An alternative reading is that the signal incorporates accredited post-serve data. The author retains the broader-scope interpretation for H1b; it is an author judgement, not an MRC finding ◆. Peer 39's acquisition post lists the accredited metrics more narrowly ◆.

Protected by Mediaocean's home page says it is MRC accredited for rendered display and video ad impressions on desktop, mobile web, in-app and CTV. MRC's letter of 22 September 2026 accredits display impressions in the first three environments and, on CTV, video impressions only, by non-SDK measurement; read literally, the claim extends display impressions to CTV ◆. The page's SIVT sentence, which names all platforms including CTV, is wider on CTV by the same reading, while the company page's narrower claim of viewability and SIVT detection stays within scope ◆ ◆.

The Zefr profile found a third claim of the same kind. Zefr's company page lists, as a 2025 milestone, MRC accreditation for viewability and IVT, while MRC accredits Zefr's YouTube video impressions and viewability with GIVT filtration only and lists Zefr under no invalid-traffic metric; the profile codes the claim broader by implication ◆.

The table separates ambiguous claims, unreadable scope and priority or uniqueness claims outside H1b's metric-scope test ◆. The verifier pass moved IAS's Quality Impression and HUMAN's MRC brief to ambiguous ◆. Whether a dated release still online counts as current remains an author decision.

PRODUCT	PAGE, AND DATE	WHAT IT SAYS	MRC'S LISTED SCOPE	CODING
Peer39 brand safety and suitability	Verification page, undated, read 10 November 2026	Viewability measurement described as MRC-accredited and built into the pre-bid signal	No entry for Peer39's product; the Peer39 Adloox service is accredited post-serve only, on desktop, mobile web and in-app (29 June 2026)	Broader on this paper's reading of the pre-bid wording, with a post-serve-input alternative ◆ ◆
Protected by Mediaocean	Home page, 2026 footer, read 10 November 2026	Accredited for rendered display and video impressions on desktop, mobile web, in-app and CTV, and for SIVT on all platforms including CTV	Display impressions on desktop, mobile web and in-app; on CTV, video impressions only, by non-SDK measurement (22 September 2026)	Broader on a literal reading ◆ ◆
Zefr brand safety and suitability	Company page, read 10 November 2026	A 2025 milestone of MRC accreditation for viewability and IVT; a 2019 milestone of industry accreditations for brand safety and suitability that names no body	YouTube video impressions and viewability with GIVT filtration only; no IVT metric listed; first Zefr decision found November 2025	Broader by implication in the Zefr profile's coding; the 2019 milestone not checkable ◆
Pixalate ad fraud protection and analytics	Blog post of 27 January 2022, live on 10 November 2026	Pre-bid blocking described as MRC-accredited	No accredited Pixalate pre-bid service; pre-bid SIVT in audit as of September 2025	Broader if a live dated post counts as current (the author's call) ◆

PRODUCT	PAGE, AND DATE	WHAT IT SAYS	MRC'S LISTED SCOPE	CODING
DoubleVerify advertiser verification	MRC one-sheet, PDF created 8 April 2026, read 10 November 2026	Pre-bid property-level brand suitability avoidance and Authentic Brand Suitability segments listed against desktop, mobile web, in-app and CTV	Pre-bid brand safety and suitability segments on desktop, mobile web and in-app; on CTV, fraud and IVT, Fully On-Screen, app-store categories and site and app lists (13 April 2026)	Ambiguous: within scope if the one-sheet's footnote 2, which lists those CTV components, governs the rows ◆
DoubleVerify advertiser verification	Fraud, brand-suitability and viewability pages, undated, read 10 November 2026	DV Authentic Ad called an MRC-accredited metric, defined to include the intended geography	Authentic Ads and Authentic Rate listed without a geography note; DoubleVerify's one-sheet says they are not accredited when campaign settings include geographic targeting	Ambiguous: within MRC's listed scope, broader than DoubleVerify's own disclosed basis ◆ ◆
DoubleVerify advertiser verification	Release of 24 April 2024, live on 10 November 2026	Accreditation extended to CTV pre-bid property-level brand suitability, contextual and Fully On-screen segments	Brand safety, suitability and contextual pre-bid segments on desktop, mobile web and in-app; app-store categories and site and app lists accredited on CTV (13 April 2026)	Ambiguous: a dated release, which may describe the app-level CTV segments ◆ ◆
IAS advertiser verification	Quality Impressions and ad-fraud pages, read 10 November 2026	Quality Impression called MRC-accredited, defined to need a suitable environment and the correct geography	Quality Ads on desktop, mobile web and in-app; the July 2026 letter excludes ad verification and lists no brand-safety or geography metric, which only the November 2025 Amazon DSP letter lists	Ambiguous ◆ ◆
HUMAN Ad Integrity Suite	MRC solution brief, PDF created 16 August 2023, still served on 10 November 2026	Calls its pre-bid accreditation one for "mitigation" of SIVT and offers end-to-end SIVT protection	Pre-bid IVT detection; MRC says detection, not prevention of ad serving (11 April 2025)	Ambiguous ◆ ◆
Pixalate ad fraud protection and analytics	Media Ratings Terminal FAQ, undated, read 10 November 2026	App and domain IVT rates from a global data pool described as MRC-accredited IVT measurement	The Analytics platform's reported statistics (8 April 2025)	Ambiguous: whether ratings derived from accredited detection are the accredited service is MRC's judgment ◆
Zefr brand safety and suitability	Brand-safety solution page, read 10 November 2026	MRC logo beside content-level brand safety and suitability and viewability, on a page presenting YouTube, Meta, TikTok, Snap and Reddit	YouTube only	Ambiguous: unqualified rather than expressly broader ◆
Adloox verification	Releases of 25 January 2023, 23 September 2024 (dateline June 2024) and 24 October 2024	The company described as MRC-accredited in announcements of Equativ pre-bid measurement, Display & Video 360 segments and Meta brand safety and suitability measurement, with no disclaimer	None of those services is accredited	Ambiguous: the descriptor attaches to the company, and the releases are dated ◆
Adloox verification	Company page, read 10 November 2026	Accredited since 2014 for comprehensive ad measurement and traffic validation, with no footnote or link to specifics	Impressions and viewability on desktop, mobile web and in-app; SIVT on desktop and mobile web	Not checkable: no metric or environment stated ◆
IAS advertiser verification	Ad-fraud page, read 10 November 2026	IAS the first verification provider accredited for SIVT on CTV	MRC's list captures of 24 January 2023 and 23 February 2024 show CTV SIVT for DoubleVerify, HUMAN, Oracle Moat, Pixalate and Protected Media before IAS's April 2024 grant	Priority claim, outside the scope test; contradicted by MRC's record ◆

PRODUCT	PAGE, AND DATE	WHAT IT SAYS	MRC'S LISTED SCOPE	CODING
Pixalate ad fraud protection and analytics	Analytics page, read 10 November 2026	Its 12 SSAI metrics called unique to Pixalate	IAS is also accredited for SSAI video metrics; MRC does not generally approve "only" claims	Uniqueness claim, outside the scope test ◆

Other screened claims omit limits without demonstrably extending scope, including certified-app restrictions and Zefr's English-only in-stream restriction ◆ ◆. Channel Factory, illuma, Fraudlogix and Escalated made no MRC claim in that screen, nor did HUMAN's June 2026 suite release ◆.

The earlier screen could read at least one page or release for each product it covered ◆. It did not establish the accreditation-claim position of mFilterIt, FouAnalytics, Pixability and Sightly. Unscreened pages cannot be coded unreadable, so the register-wide unresolved gate has not been evaluated completely. Retaining H1b's contradicted verdict requires accepting the stated scope interpretations and a counterexample-first reading of the rule.

How H1a counts units

H1a predicted that at least three quarters of the units the scored products sell hold current MRC accreditation. A unit is a product, a metric family and an environment. The six families are display viewability, video viewability, GIVT, SIVT, property-level ad verification, and content-level brand safety or suitability; the environments are desktop web, mobile web, in-app, CTV, audio and digital out-of-home, plus each large platform on which a metric is sold. Attention metrics belong to the attention paper and are not units here. What a product sells comes from its own documentation, recorded in the product register; what is accredited comes from MRC's accredited list, under-review list and letters at the freeze, read for product, metric, environment, format, placement, language and standard. A letter's placement and language exclusions narrow what an accredited unit covers: Zefr's English-only in-stream grant does not reach the Shorts, Google Video Partner inventory or other languages Zefr sells on YouTube ◆.

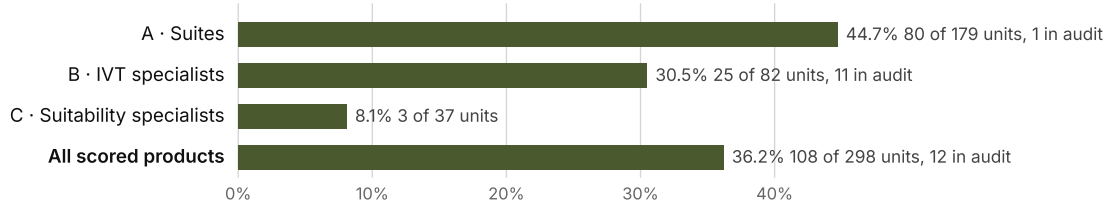
Under-review units count as unaccredited, and IAS's property-level row has no current letter basis ◆ ◆. H1a reports pooled, cohort, environment and content-level shares with both counts. Its registered support rule requires at least 75% overall and no cohort below 50%. It is contradicted below either threshold and unresolved if sold units cannot be established for more than a quarter of scored products.

The grid matters as much as the list. In a preliminary count during the research, whether Zefr's YouTube viewability reporting counted as one unit or three, with or without GIVT units, moved the suitability cohort's share across the 50% floor ◆. One grid has to apply to every product, and the units come from the scored register.

The merged documented grid records 108 accredited units out of 298 sold (36.2%): cohort A 80/179 (44.7%), cohort B 25/82 (30.5%) and cohort C 3/37 (8.1%). Content-level brand safety or suitability accounts for 1 accredited unit out of 61 sold units in that family; 1 further content-level unit is under review. Across all families, 12 units are under review (1 in A, 11 in B and none in C), counted as unaccredited. H1a is contradicted: the pooled share is below the 75% rule and every cohort is below its 50% floor. These are documented product×metric-family×environment units, not shares of impressions or spend.

illumaSentinel has an unknown sold-unit denominator and remains outside this documented tally; unknown is not zero, and those omissions do not meet the registered more-than-a-quarter unresolved gate. Status rests on the November records, pending the 11 December 2026 freeze recheck ◆. One counting ruling matters here: MRC's accredited list still shows IAS's property-level verification on desktop and mobile web, while the July 2026 letter linked from it omits them. The grid follows the list, as it does for every product, and counts both units as accredited; excluding them gives 106 of 298 (35.6%). Neither reading changes IAS's accreditation score or the H1a verdict ◆.

MRC-ACCREDITED SHARE OF THE UNITS EACH COHORT SELLS



A unit is a product x metric family x environment that the product documents as sold, narrowed by the letters' placement and language exclusions. Units in audit count as not accredited.

> **Under half of what the scored products sell is MRC-accredited.** Accredited units divided by sold units, for each cohort and pooled, from the cell-verified scoring of each product (hypothesis H1a). A sold unit is counted only where the product's own documents or a verified record show it is sold.
>
> Units: percent of sold units · Scope: the scored products at 10 November 2026 · Sources [47]
>
> The denominator depends on what each vendor documents as sold; an undocumented product line is not counted.

CHAPTER 3

What the independent record shows

In the sources screened to 10 November 2026, no independent test reports an error rate for a scored product on a labelled set with a published denominator, supporting H3 ◆. The record measures prevalence, vendor disagreement and failure cases ◆. Adalytics' candidate miss rates concern IAS's unscored publisher tool and give no session counts ◆. Neither those cases nor vendors' replies supply the test H3 requires ◆.

IN SHORT

The ANA's 2023 study measured IVT at 0.78% of impressions and MFA at 21%, using vendors' and a data supplier's classifications ◆ ◆. DoubleVerify and IAS announced their MFA products after that study's window ◆ ◆. Adalytics' ten reports naming a verifier are observational and publish no product error-rate denominator ◆ ◆. Smith et al. found vendor disagreement on news articles, with labels partly constructed by the authors and no ground truth ◆. The candidate bot miss rates concern IAS Publisher Optimization, which this paper does not score, and give no session counts ◆ ◆. H3 is supported on this dated, bounded screen ◆. Vendors' protected-versus-unprotected comparisons use their own classifiers and publish no arm sizes ◆. Industry ad-fraud totals are sponsored projections or models with incompatible definitions ◆.

What the ANA's 2023 study measured

The ANA's study measured what the participants' vendors and one data supplier found in 21 advertisers' buys; it tested no product. The Programmatic Media Supply Chain Transparency Study analysed log-level data from 21 ANA members, of 67 that wanted to take part, covering $123 million of open-web programmatic spend and 35.5 billion impressions from September 2022 to January 2023 ◆. Three demand-side platforms (DSPs), six supply-side platforms (SSPs) and three verification companies, DoubleVerify, IAS and Oracle Moat, supplied data and had no role in the analysis (p. 17); the panel selected itself ◆. The Complete Report (December 2023) is cited by printed page from a public copy hosted by Adslot [69]. The logs were reconciled by TAG TrustNet, a joint initiative of Fiducia and TAG, the Trustworthy Accountability Group, which certifies verification vendors ◆ ◆.

The participating vendors' classifications were matched to DSP logs ◆. IVT was 0.78% of impressions and 0.69% of spend, median measurability was 86%, and 85% of measurable impressions were viewable ◆ ◆. These are measured prevalence, and the ANA's attribution of low IVT to protection is its interpretation ◆.

Domains matched to DeepSee.io's 4,500-site MFA list took 21% of impressions and 15% of spend in the subset analysed [◆]. The ANA reported good viewability and low IVT and described the sites as usually brand safe [◆] [◆]. Yet brand safety was outside the study, and the Q3 2025 benchmark cites inconsistent vendor methods as the reason [◆].

The participating DoubleVerify and IAS services announced MFA products after the study window, in September and October 2023 respectively [◆] [◆]. DoubleVerify's earlier on-request handling remains undated, and no participant contract establishing it in force was found in the screened sources [◆]. Adloox, which supplied no study data, did advertise an MFA-labelled pre-bid segment by December 2022, but equivalence to the September 2023 joint definition is undocumented [◆] [◆]. The ANA advised testing the later offerings against what they promised [◆].

The benchmark that followed

The ANA's quarterly Programmatic Transparency Benchmark repeats the design at a larger scale and keeps its limit: it reports prevalence as its suppliers measure it and tests none of them [◆]. IVT, measurability and viewability come from log data matched between DSPs and the advertisers' own verification providers, MFA data from DeepSee.io and, from the third quarter of 2025, brand-risk data from Mobian, a contextual vendor; the panel changes every quarter [◆] [◆].

RELEASE	PERIOD	MARKETERS (PARTICIPATING / ACTIVE)	IMPRESSIONS	SPEND	IVT	MFA	EVIDENCE
2023 study (Complete Report)	Sep 2022 to Jan 2023	21	35.5 billion	$123 million	0.78% of impressions, 0.69% of spend	21% of impressions, 15% of spend (subset)	[◆] [◆] [◆]
2024 Benchmark, Findings Update	Apr to Sep 2024	37 / 21	38.5 billion	$235 million	not in the passages read	median 1.1% (10% in the 2023 study)	[◆] [◆]
Q3 2025	Jul to Sep 2025	39 / 21	17.5 billion	$142 million	0.4% of spend (waterfall)	0.5% of spend (waterfall; 0.4% in the summary)	[◆] [70, 71]
Q4 2025	Oct to Dec 2025	54 / 35	16.7 billion	$142 million	0.4% of spend	0.4% of spend	[◆] [72, 71]
Q1 2026	Jan to Mar 2026	86 / 66	20.9 billion	$160 million	0.2% of spend	1.1% of spend	[◆] [◆]
Q2 2026 (release read in Fiducia's repost; the report is behind a form)	Apr to Jun 2026	not read	not read	not read	spend lost "held at 0.1%"	not read	[◆] [73]

The benchmark's changing panel and supplier classifications cannot separate cleaner supply from improved detection [◆]. It warns that waterfall values and medians differ, and the Q2 2026 release's description of 0.1% IVT as unchanged conflicts with Q1's printed 0.2% [◆]. Mobian's separate claim that legacy systems misclassified roughly a third of professional news pages gives no labels or denominator [◆]. A July 2026 ANA, TAG and Fiducia analysis found AI-generated pages passing the screened controls but named no verifier, leaving H3 unchanged [74] [◆].

Adalytics' reports, case by case

Adalytics' reports document failure modes: cases, without denominators, of ads carrying a verifier's code that ran where buyers would not want them or reached non-human traffic. From June 2023 to March 2025 it published ten reports that name a verification vendor, all observational, and the three that bear most on DoubleVerify, IAS and HUMAN make no claim about magnitude or relative abundance ◆. Adalytics sells analytics to advertisers and declares that no publisher or ad-tech vendor funds its research ◆.

REPORT (SUBJECT)	DATED, AND ON WHAT BASIS	VERIFIERS NAMED, AND THEIR ROLE	METHOD AND SAMPLE	PERIOD OBSERVED	ESTIMAND	EVIDENCE
Google TrueView and Google Video Partners	Public by 27 June 2023 (Google's reply that day; first archived 28 June)	IAS and DoubleVerify, said to rely on YouTube's data feeds; MRC	Case report; no sample stated	Not stated	Failure-mode cases	◆
MRC's answers on TrueView	July 2023 (sitemap 14 July; first archived 23 July)	MRC; Oracle Moat, DoubleVerify, IAS	Relays MRC's statement that no third party was audited or accredited for Google brand safety	2023	Record of accreditation scope	◆
YouTube advertisers and children's data	17 August 2023 (sitemap and first archive)	IAS and DoubleVerify in YouTube's account of beta tests, IAS also as a recipient of one landing page's data; Channel Factory once, as a vendor Apple appeared to use	Case report; non-exhaustive list of brands	June and July 2023	Cases, not about verifiers' classifications	◆ [75]
Search ads on objectionable pages	November 2023 (first archived 28 November; sitemap 6 December)	HUMAN and DoubleVerify as advertisers whose search ads appeared on such pages	Case report	2023	Cases, not about verifiers	[76]
MFA sites in January 2024	March 2024 (sitemap 11 March; first archived 18 March)	IAS, DoubleVerify, Oracle Moat, Pixalate, HUMAN, Fou Analytics: tags in ads on MFA sites	22 sites judged MFA; panel, crawlers and archives; brands counted, no impression denominator	January 2024	Failure-mode cases	◆ ◆
Forbes www3 subdomain	Undated; public by 4 April 2024 (first archived that day); Forbes made contact on 20 March, the page-URL mis-declaration was no longer observable from 22 March and the subdomain was offline from 2 April	IAS, DoubleVerify, Oracle Moat, HUMAN tags in ads; their publisher code absent	One publisher; mis-declared bid requests to four SSPs	To 22 March 2024	One failure-mode case	◆
User IDs in ad auctions	10 May 2024 (sitemap and first archive)	HUMAN and Oracle Moat as an SSP's declared partners; MRC	One case	2024	Case, not about verifiers' classifications	◆ [77]

REPORT (SUBJECT)	DATED, AND ON WHAT BASIS	VERIFIERS NAMED, AND THEIR ROLE	METHOD AND SAMPLE	PERIOD OBSERVED	ESTIMAND	EVIDENCE
"AI" brand safety technology	7 August 2024 (first archive; same-day replies)	IAS publisher key-values; DoubleVerify and IAS tags; MRC	Convenience sample	Explicit-page examples undated; dated examples April 2023 to July 2024	Failure-mode cases, both directions	◆
Image host imgbb.com and ibb.co	7 February 2025 (first archived that day); the senators' letters cite "Jan. 2025", its data cutoff	DoubleVerify and IAS tags on explicit adult pages; MRC	Thousands of URLScan.io recordings and its own sessions; advertisers' domain-level account relayed	2017 to January 2025	Failure-mode cases and a second-hand domain-level account	◆ ◆
Pre-bid bot filtration	28 March 2025 (sitemap; dated replies)	DoubleVerify, IAS, HUMAN, MRC	Three crawl datasets: over a petabyte, over two million websites, over seven years; no counts published	2017 to March 2025	Failure-mode cases; one candidate rate (IAS flag)	◆ ◆

The explicit-pages report of 7 August 2024 began when a multinational brand, told that its campaigns were "100% brand safe", asked for a review; the phrase is this report's ◆. It showed IAS's publisher-side endpoint returning "veryLow" or "low" on all seven brand-safety keys, as IAS's own guide defines them, for explicit pages, mostly Fandom wikis, and ads carrying DoubleVerify or IAS code there; the same endpoint labelled mainstream news "medium" or "high" ◆. The sample was one of convenience, and the report disclaims any claim about a classifier's sensitivity, specificity or accuracy ◆. Its key-values come from a publisher product, not the IAS advertiser verification this paper scores ◆.

The image-host report of 7 February 2025 works at the level of a domain. It reviewed thousands of URLScan.io recordings of imgbb.com and ibb.co from 2017 to 2024 and relayed several advertisers' statements that their vendors had marked 100% of measured impressions there safe or suitable, a second-hand, domain-level account rather than a measured rate ◆. The ads it showed, some with DoubleVerify or IAS tags, sat beside adult explicit pages; its one archived instance involving putative child sexual abuse material (CSAM) was reported to the authorities and not shown, and advertisers could not count ads beside CSAM for lack of page-level reporting ◆. A domain that hosts illegal images is not the same as an ad beside one.

The pre-bid bot report of 28 March 2025 shows ads served to bots in three crawl datasets, URLScan.io, HTTP Archive and one it does not name ◆. Its statement that advertisers were billed rests on non-public log files and invoices from its work for Fortune 500 advertisers, and it disclaims any finding of fault, magnitude or prevalence ◆. Whether URLScan's crawler declares itself is disputed: DoubleVerify and MRC say its user agent does not, while HTTP Archive is on the Spiders and Bots list of the Interactive Advertising Bureau (IAB) ◆ ◆. Senator Warner's letters, dated 24 March 2025 and released on 28 March, asked the Federal Trade Commission (FTC) and the Justice Department to investigate DoubleVerify, IAS and HUMAN, and footnote 6 of the FTC letter describes the failures as general invalid traffic (GIVT), not advertising fraud under industry definitions; the letters record questions, not findings ◆. The report's HUMAN cases concern platforms that, MRC says, make their own bid decisions on HUMAN's signals, and no HUMAN response was found by 10 November 2026 ◆. DoubleVerify sued Adalytics over this report on 12 May 2025, and the

pleading-stage ruling of 27 April 2026 took DoubleVerify's allegations as true, made no finding on the report's accuracy and left the case to continue (Chapter 9) ◆.

One figure is framed as a detection error rate for a named product. Treating its crawls as "100% confirmed bot traffic", the report says the "fr" flag of IAS's Publisher Optimization marked the declared HTTP Archive bot valid 16% of the time in 2022 to 2024 crawls (17% in its summary) and URLScan.io's bot 77% of the time in 2019 to 2024 crawls of over 1,200 websites, without publishing the number of sessions behind either share ◆. IAS's own implementation guide of 11 July 2024 defines the flag as "Whether the ad represents invalid traffic" [78] ◆. MRC said on 11 April 2025 that IAS's publisher tools were neither audited nor accredited, and its list read on 10 November 2026 still shows none, though it now lists IAS's pre-bid IVT detection for desktop, mobile web and in-app, accredited from June 2026 ◆ ◆.

What the vendors and MRC answered

Every answer corrects something, and none publishes the labelled data that would settle the case it answers. H3 judges a response on its merits where it corrects scope, data, decoding, staleness, the denominator or the billing stage, without requiring a new test of equal size ◆.

DATE	RESPONDENT	REPORT ANSWERED	WHAT IT CORRECTS	DATA PUBLISHED	EVIDENCE
27 June 2023, updated 13 July	Google	Google Video Partners	Method; cites DoubleVerify's and IAS's 93% viewability on Google Video Partners as contradiction	No counts	◆
12 July 2023	IAS	Google Video Partners	Data: 3.4% of its YouTube and partner impressions, 92.74% viewable, 0.40% IVT (January to early July 2023)	Aggregate shares	◆
7 August 2024	DoubleVerify	Explicit pages; also the MFA and Forbes reports	Scope (all tags shown to it were publisher services), settings (exception lists, monitoring); says all content shared was classified accurately, virtually all the MFA sites were already classified as MFA and the Forbes inventory as low-tier MFA	No classifications or counts	◆ ◆ ◆
7 August 2024	IAS	Explicit pages	Scope ("almost all" tags from publisher solutions), settings (monitoring, not blocking); no timeframe or sample in the report; silent on the endpoint's low-risk values	No	◆
10 January 2025	DoubleVerify	Bot report, in draft	Billing stage: under the industry standard, GIVT that reaches an ad is removed from reporting and not charged	No	◆
7 February 2025	DoubleVerify	Image host	Denominator: 0.000047% of the media transactions it measured in 30 days; the archived copy adds that 83% served to the home page, a figure the live page dropped on 15 October 2025; classification only above an ad-traffic threshold	One share	◆ ◆
7 February 2025	IAS	Image host	Denominator: 88,000 of nearly 1.9 billion webpages in 60 days; domains made ineligible for monetization	Page counts	◆
13 February 2025	DoubleVerify	Image host (remedial)	New default-on "Highly Illicit: Do Not Monetize" category; about 100 of over 300 domains in public data from the National Center for Missing & Exploited Children (NCMEC) already covered	Domain counts	◆
14 February 2025	IAS	Image host (remedial)	Over 350 NCMEC-notified sites already covered; a category for unmoderated user-generated content sold since 2023	Site and campaign counts	[79]

DATE	RESPONDENT	REPORT ANSWERED	WHAT IT CORRECTS	DATA PUBLISHED	EVIDENCE
21 March 2025	IAS	Bot report, before publication	Method: publishers can target away from impressions labelled IVT	No	[80]
28 March 2025	DoubleVerify	Bot report	Every example shared before publication correctly identified; billing stage	No	
28 March and 3 April 2025	IAS	Bot report	Scope (publisher tag, not advertiser tag), stage (no post-bid blocking by design), decoding (back-end adjudication); HTTP Archive flagged as GIVT, URLScan.io impressions as GIVT, sophisticated invalid traffic (SIVT) or reduced-value inventory	No counts	
7 April 2025	DoubleVerify	Bot report	Scope: of 115 examples, 62 had no DoubleVerify measurement, 10 fell outside retention, all 43 matched were invalidated post-bid; labelling (URLScan undeclared); decoding (the "DV source code" is DSP code); its own test figures	Example counts; no per-impression data or method	
11 April 2025	MRC	Bot report	Scope (IAS publisher tools and pre-bid IVT, DoubleVerify publisher tools and Scibids then unaudited); standard (post-serve GIVT filtration required, pre-bid not); labelling (URLScan.io's user agent undeclared); no insight into billing	Did not test each observation	
None found by 10 November 2026	HUMAN	Bot report	No public response in the searches and sitemaps read	n/a	

The August 2024 answers correct scope and settings, but not the low-risk values IAS's publisher endpoint returned for explicit pages, which are classifications whatever a buyer's settings . In February 2025 both vendors supplied the denominators the report lacked, and DoubleVerify disclosed a coverage limit: a page is classified only once it reaches a certain level of ad traffic . Its statement as archived on 7 February 2025 adds, in its summary, "with over 80% served to the site's home page", and puts that share at 83% in the body; the live page, revised on 15 October 2025, carries neither [81] . Neither vendor met the advertisers' account of 100% safe or suitable with data .

On the bot report, DoubleVerify's breakdown of the 115 examples is a scope correction H3 counts, since more than half carried no DoubleVerify measurement; its own detection figures come from the classifier being defended, without method or sample . IAS corrects scope, stage and decoding by its own assertion, and MRC supports it on scope and stage, but IAS classified the URLScan.io impressions as GIVT, SIVT or reduced-value inventory, a category it keeps apart from IVT, so its answer does not say every one was classified as invalid . MRC, which had not tested each observation, wrote: "It's very likely a material amount of SIVT goes undetected" [82].

Smith et al. 2026: agreement, not accuracy

Smith, Grossman, Torres-Agüero, Sen, Borcea and Chen, five authors at the New Jersey Institute of Technology and one at DeepSee.io, measured how often three providers agree about the same news article, not whether any is right . The paper was posted to arXiv on 3 January 2026 and published on 25 May 2026 in the proceedings of the International AAAI Conference on Web and Social Media, with funding from two National Science Foundation

grants, the Martin Tuchman '62 Chair Endowment and the Leir Foundation ◆ ◆. DeepSee.io also supplied the ANA study's MFA list, and this paper records the two as one interest [69].

The authors read the brand-safety key-values that publishers' own Google Ad Manager integrations exposed, client-side, for DoubleVerify, IAS and Oracle (Moat) on news articles collected in September 2024 ◆. These are publisher-side labels, not scores obtained from the providers or the ratings advertisers buy, and the sample covers only pages whose publishers collect them ◆. Of 5,241 URLs attempted, 4,352 returned tag data and 3,219 carried any classification; 1,416 of the 3,219 (43.99%) were unsafe for at least one provider under the authors' own rules ◆. Those rules are partly the authors' construction: IAS gives no overall rating, so an article counted as unsafe for IAS when any flag was present or any category sat above "veryLow", and DoubleVerify's label is read from its "Any moderate content" code ◆ ◆. A union across three providers is not a ground-truth unsafe rate. The data date from the month Oracle retired Moat, on 30 September 2024, so the Oracle comparisons are historical ◆ ◆.

PROVIDERS COMPARED	ARTICLES	DISAGREEMENT	KRIPPENDORFF'S ALPHA	EVIDENCE
IAS and DoubleVerify	94, all on forbes.com (89 also rated by Oracle)	37.23%	0.19	◆ ◆ [83]
Oracle and IAS	329, rated by just these two	29.79%	0.42	◆ [83]
Oracle and DoubleVerify	249, rated by just these two	14.06%	0.53	[83]
All three	89, all on forbes.com	41.57%	0.35	

The overlaps are small, and the IAS and DoubleVerify pair rests on 94 articles from one publisher ◆. Of the 31 articles Oracle or IAS flagged for hate speech, the two disagreed on 27 ◆. The authors found misclassifications by every provider by hand, without a rate, say no ground-truth dataset of brand-safety labels exists, and call for a public, human-labelled benchmark ◆. The study measures consistency between vendors, not accuracy.

Other academic work and platform reports

Few other studies touch a verification product. Adalytics' analysis of December 2020, inside H3's window from 2019, found Comscore and Oracle Moat agreeing on safe or unsafe for 58.7% of 25,730 archived wsj.com articles; a companion post was edited after IAS said the data it described were not advertiser keyword blocklists ◆. Searches of arXiv on 10 November 2026 found no other study testing a named verification product's outputs and none reporting a verification product's miss or false-positive rate on a labelled set ◆. The labelled designs found test other products, such as FP-Inconsistent's evasion rates against the bot detectors DataDome and BotD, while Vekaria and colleagues document inventory pooling that evades brand-safety protections ◆. Google Scholar, SSRN, the ACM Digital Library and IEEE Xplore were not queried ◆. Platform transparency reports, YouTube's Violative View Rate and Meta's prevalence metric, estimate violating content from labelled samples, but they are first-party measures of each platform's own policy enforcement, not error rates of a third-party product ◆.

Independent and platform studies, by what they can estimate

Each study as its design allows: failure-mode cases, prevalence in a sample, agreement between vendors, a candidate error rate, or a platform's own prevalence metric. Vendor and commissioned evidence is kept out of this table and discussed apart in chapter 3. Sponsorship is from the source register.

Full records with limitations and source passages are in data/study-evidence.csv.

STUDY	YEAR	DESIGN	SAMPLE AND DENOMINATOR	PRODUCTS NAMED	CATEGORY	ESTIMAND IT SUPPORTS	RESPONSES, AND WHAT THEY CORRECT	WAS THE DETECTION SOLD THEN?	SPONSOR	OBSERVATION PERIOD	SRC
ANA Programmatic Media Supply Chain Transparency Study (First Look and Complete Report)	2023	reconciled log-level data, self-selected panel	21 advertisers, $123m, 35.5bn impressions; MFA on a subset	DoubleVerify, IAS, Oracle Moat (data suppliers)	IVT, viewability, measurability, MFA	prevalence in a sample, measured by the products themselves; MFA by DeepSee.io list	none from vendors found	IVT and viewability yes; MFA no (DV beta 20 Sep 2023, IAS beta 16 Oct 2023)	Commissioned and published by the ANA (advertisers' trade association). Data processed by TAG TrustNet, a unit of TAG, which certifies...	Sep 2022 to Jan 2023	[69] [154] [155] [156]

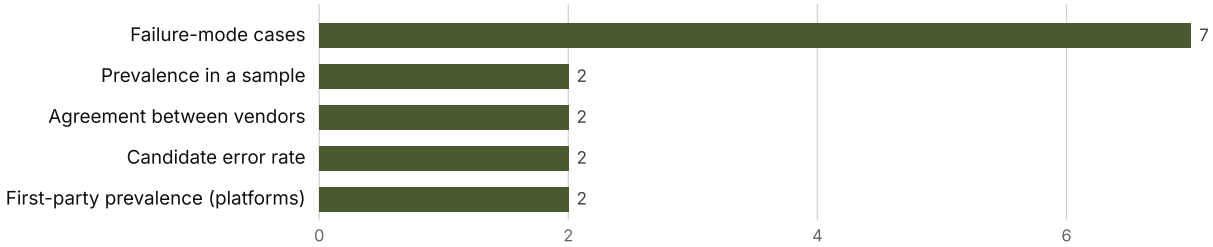
STUDY	YEAR	DESIGN	SAMPLE AND DENOMINATOR	PRODUCTS NAMED	CATEGORY	ESTIMAND IT SUPPORTS	RESPONSES, AND WHAT THEY CORRECT	WAS THE DETECTION SOLD THEN?	SPONSOR	OBSERVATION PERIOD	SRC
ANA Programmatic Transparency Benchmark (with TAG TrustNet, Fiducia)	2024	quarterly panel, cost waterfall and medians	from 21 advertisers (2024) to 86 participating / 66 active (Q1 2026)	verification providers unnamed in method; DeepSee.io; Mobian	IVT, viewability, measurability, MFA, brand risk	prevalence in a panel, measured by suppliers	none found	yes for IVT and viewability; MFA sold from late 2023	ANA with TAG TrustNet and Fiducia. TAG certifies verification vendors; Fiducia runs the data exchange and markets participation in the...	Apr 2024 to Jun 2026	[157] [158] [159] [70] [72] [71] [73]
Adalytics, Google Video Partners / TrueView	publ	observational case report	none stated	IAS, DV (YouTube Measurement Program)	brand safety, viewability on GVP	failure-mode cases	Google 2023-06-27 and 2023-07-13; IAS 2023-07-12: data (vendor-measured prevalence); not a test either way	yes (measured via platform data, per report)	Adalytics Research LLC: sells data science and analytics services to advertisers; its About page self-declares that it does no research on...	not stated	[160] [161] [162] [163]
Adalytics, MFA sites in January 2024	firs	observational case report	22 sites; brands counted, no impression denominator	IAS, DV, Oracle Moat, Pixalate, HUMAN, Fou Analytics (tags)	MFA	failure-mode cases	DV 2024-08-07: settings (classification vs avoidance), unverified	beta (DV all advertisers; IAS select customers)	Adalytics Research LLC: sells data science and analytics services to advertisers; its About page self-declares that it does no research on...	Jan 2024	[164] [165] [156]
Adalytics, Forbes www3 subdomain	publ	observational case report	one publisher	IAS, DV, Oracle Moat, HUMAN (tags)	MFA, domain misdeclaration (SIVT type)	failure-mode case	DV 2024-08-07: classification claim, unverified	MFA beta/GA; SIVT yes	Adalytics Research LLC: sells data science and analytics services to advertisers; its About page self-declares that it does no research on...	to 22 Mar 2024	[166] [165]
Adalytics, "AI" brand safety technology	2024	observational case report, convenience sample	none	IAS (publisher key-values), DV and IAS (tags)	brand safety	failure-mode cases	DV and IAS 2024-08-07: scope (publisher products), settings; no data published	yes	Adalytics Research LLC: sells data science and analytics services to advertisers; its About page self-declares that it does no research on...	dated examples Aug 2023 to May 2024	[167] [165] [168]
Adalytics, image host imgbb.com / ibb.co	2025	observational case report from crawler archives	thousands of recordings; no denominator; 100% is advertisers' domain-level account	DV, IAS (tags)	brand safety (adult); CSAM-hosting domains	failure-mode cases	DV and IAS 2025-02-07; DV 2025-02-13: denominator (share of vendor volume), stage (measurement vs blocking); not the 100% account	adult category yes; CSAM-domain category no (DV launched 2025-02-13)	Adalytics Research LLC: sells data science and analytics services to advertisers; its About page self-declares that it does no research on...	2017 to Jan 2025	[169] [170] [171] [172]

STUDY	YEAR	DESIGN	SAMPLE AND DENOMINATOR	PRODUCTS NAMED	CATEGORY	ESTIMAND IT SUPPORTS	RESPONSES, AND WHAT THEY CORRECT	WAS THE DETECTION SOLD THEN?	SPONSOR	OBSERVATION PERIOD	SRC
Adalytics, pre-bid bot filtration (ads served to bots)	2025	observational, three crawler datasets	no counts published	DV, IAS, HUMAN; DSPs and SSPs	GIVT, pre-bid bot avoidance	failure-mode cases	DV 2025-01-10, 2025-03-28, 2025-04-07; IAS 2025-03-28, 2025-04-03; MRC 2025-04-11; HUMAN none found: billing stage, labelling (URLScan undeclared), decoding (DSP code); DV publishes no counts	yes	Adalytics Research LLC: sells data science and analytics services to advertisers; its About page self-declares that it does no research on...	2017 to 2025	[173] [174] [175] [176] [177] [178] [82]
Adalytics, pre-bid bot filtration (IAS publisher "fr" flag)	2025	all-bot labelled crawls	shares only (16% or 17% declared; 77% URLScan.io); session counts not published	IAS Publisher Optimization	IVT flag (publisher side)	candidate error rate (miss rate), denominator unpublished	IAS 2025-03-28 and 2025-04-03; MRC 2025-04-11: scope, stage and decoding (IAS assertion; MRC confirms scope)	yes (sold to publishers); not MRC-accredited; not a scored row on the seed product register	Adalytics Research LLC: sells data science and analytics services to advertisers; its About page self-declares that it does no research on...	2019 to 2024	[173] [177] [178] [82]
Adalytics, news articles labelled unsafe	2020	archive extraction of publisher-side labels	25,730 wsj.com articles	Comscore, Oracle Moat (also Grapeshot, IAS)	brand safety	agreement between vendors (58.7%)	IAS (relayed on the companion post): decoding (keyword lists); Adalytics edited	yes	Adalytics Research LLC: sells data science and analytics services to advertisers; its About page self-declares that it does no research on...	to Dec 2020	[179] [180]
Smith et al., ICWSM 2026	arXi	cross-sectional audit of publisher-side key-values	3,219 classified of 4,352 with tag data of 5,241 attempted	DoubleVerify, IAS, Oracle (Moat)	brand safety	agreement between vendors	none found	yes (Moat closing that month)	NSF grants CNS 2237328 and DGE 2043104, the Martin Tuchman '62 Chair Endowment and the Leir Foundation; one author from DeepSee.io...	Sep 2024	[83] [181] [182]
Venugopalan et al., FP-Inconsistent	arXi	honey site with purchased bot traffic (labelled)	about 500,000 requests from 20 bot services	DataDome, BotD (not ad verification)	bot detection	error rate (evasion) for non-verification products	n/a	n/a (not scored products)	not read (abstract only)	not stated in abstract	[183]
Vekaria, Nithyanand, Shafiq	arXi	supply-chain measurement	not read (abstract only)	ad exchanges (no verifier named in abstract)	brand safety circumvention	failure mode (dark pooling)	n/a	n/a	not read (abstract only)	not stated	[184]
YouTube Violative View Rate	2021	labelled sample of videos (first-party)	not stated in post	YouTube (own enforcement)	policy-violating content	prevalence, first-party	n/a	n/a (first-party)	Google/YouTube describing its own enforcement metric (first-party).	Q4 2020	[185]
Meta prevalence metric	page	labelled samples of content views (first-party)	not stated	Meta (own enforcement)	policy-violating content	prevalence, first-party	n/a	n/a (first-party)	Meta describing its own enforcement metric (first-party).	not stated	[186]

Coding the record by estimand

Coded by the estimand each design supports, the record is mostly failure cases. Workstream B coded 15 rows: 7 failure-mode cases, 2 studies of prevalence in a sample, 2 of agreement between vendors, 2 candidate error rates and 2 platforms' first-party prevalence metrics ◆. The bot report fills two rows, its served-ad cases as failure modes and its IAS shares as a candidate rate; the other candidate, FP-Inconsistent, tests bot detectors that are not verification products ◆. The verifier pass added the July 2026 AI-slop analysis to prevalence and Smith et al.'s hand-found misclassifications to the failure cases ◆.

INDEPENDENT AND PLATFORM STUDIES BY WHAT THEY CAN ESTIMATE



Of the two candidate error rates, one concerns IAS's publisher-side invalid-traffic flag on declared bots in one crawl dataset (2025) and the other a non-verification fingerprinting tool. No labelled test of a scored product's sold category was found.

> **Most of the independent record shows failure modes, not error rates.** The 15 independent and platform studies coded in workstream B by the estimand each design supports: failure-mode cases without a denominator, prevalence in a sample, agreement between vendors, a candidate error rate with a labelled set and a denominator, and platforms' own prevalence metrics.
>
> Units: count of studies · Scope: studies concerning verification products, 2019 to 10 November 2026 · Sources [69] [154] [155] [156] [157] [158] [159] [70] [72] [71] [73] [160]
>
> A count of designs, not of verdicts; one study can support only one estimand here.

H3's registered support rule requires no independent error-rate study of a scored sold category on a labelled set with a denominator ◆. The screen through 10 November 2026 meets it because IAS Publisher Optimization is outside the scored register ◆. Adding that publisher tool before its session counts are public would invoke the unresolved branch for an incompletely documented candidate study ◆. The screen must be repeated at the freeze.

An independent miss-rate, false-positive or precision test of a scored sold category, with labels and a denominator, would contradict H3 ◆. The IAS candidate would matter if its session counts were published and the publisher tool entered the register ◆ ◆. Smith et al. call for a public labelled benchmark, which offers another route ◆.

Vendor and commissioned evidence, labelled as such

Vendor and commissioned numbers measure prevalence, compare protected with unprotected campaigns, including in connected TV (CTV), or model savings. In the publications screened, by or with DoubleVerify, IAS, HUMAN, Pixalate, Zefr, Peer39 and Channel Factory, plus studies by TAG, Juniper Research, CHEQ, the ANA with White Ops

and Forrester, none reports a miss rate, false-positive rate or precision for a sold category on a labelled holdout with a denominator; Adloox, Protected by Mediaocean (formerly Protected Media), Fraudlogix, Escalated, TrafficGuard, Pixability and Apostra were not screened for such claims in that workstream, and illuma and Proximic in one query only ◆.

VENDOR OR BODY (SPONSOR)	FIGURE	BASE AND PERIOD	DESIGN	EVIDENCE
DoubleVerify, release of 7 May 2026	CTV fraud under 1% protected, nearly 9% unprotected	"Billions of impressions" and controlled tests; no period, counts or assignment	Protected against unprotected, measured by DoubleVerify	◆
IAS, 20th Media Quality Report	Fraud 10.9% non-optimized, 0.7% optimized	Second half of 2024; non-optimized campaigns 4.0% of those monitored (2024 in the text, the second half on the chart)	Self-selected arms, measured by IAS	◆
IAS, 21st Media Quality Report	CTV IVT 9.1% non-optimized, 0.1% optimized	Second quarter of 2026; arm sizes not published	Same	◆
IAS, 21st Media Quality Report	IVT 1.1%	2025; open-web desktop and mobile web; IAS's intent-based definition	Prevalence	◆
Pixalate, releases of 29 April and 12 August 2026	IVT 20% web, 39% mobile app, 25% CTV (first quarter); 20%, 41% and 26% (second quarter)	Over 82 billion and over 80 billion impressions, predominantly open auction; Pixalate labels its first-quarter figures opinions	Prevalence	◆
TAG, conducted by The 614 Group	IVT under 1% in TAG Certified Channels	January to June 2024; 745 billion certified-channel impressions from agency holding companies, measured by vendors TAG describes as MRC-accredited; US	Prevalence after filtering	◆

DoubleVerify and IAS publish the same design: fraud in protected campaigns set against unprotected ones, measured by the classifier that did the protecting, with no campaign or impression counts per arm ◆. In IAS's comparison buyers chose their arm, and only 4.0% of the campaigns it monitored ran unprotected ◆. A low protected rate shows that the controls and the classifier agree; it cannot show what both missed ◆. The prevalence figures move with base and definition, so global rates from 0.1% to 41% sit in the table without contradicting one another; none is an error rate, and they cannot be pooled ◆ ◆. IAS's brand-safety fail rates appear only as multiples of baselines it does not publish ◆.

Vendors also state accuracy without a test. IAS says it is "42% more accurate than the next best provider at classifying online content", and its FY2024 10-K claims "up to 74% more accurate" brand-suitability reporting on social platforms, neither with method, sample, period or denominator ◆. HUMAN says its brand-safety engine is benchmarked against over 400,000 expert-labelled pages with "high accuracy", and gives no figure ◆. DoubleVerify's FY2025 10-K concedes that partners' data have caused errors in its technology, without a rate ◆.

Zefr's August 2025 study tests six general-purpose multimodal models on 1,500 videos, with the best F1 score at 0.91 against reviewers' labels ◆ ◆. It does not test Zefr's sold classifier ◆. The Oracle-commissioned Forrester work combines two economic-impact studies for since-retired products ◆ ◆. Zefr's Forrester survey and Channel Factory's MAGNA study measure opinions and adjacency effects ◆.

Scale claims use different units and describe populations, not error-rate denominators: DoubleVerify files about 9.5 trillion Media Transactions Measured for 2025, its billing unit,

IAS cites over 300 billion "digital interactions" a day, and Pixalate over 82 billion impressions in the first quarter of 2026 ◆ ◆ ◆.

Industry "ad fraud" totals are sponsored estimates

The industry-wide "ad fraud" totals in circulation are projections or models, each produced with or distributed by a fraud-protection vendor, or by bodies scoring their own programmes; none is a measured market total ◆ ◆.

ESTIMATE	VALUE	YEAR	BASE (DENOMINATOR)	PRODUCER AND INTEREST	STATUS	EVIDENCE
Juniper Research, ad-fraud whitepaper (September 2023)	$84.2 billion; almost $100 billion; $172.3 billion	2023; end of 2024; 2028	Global; 2023 put at 22% of digital ad spend; no method stated; a broad and a click-based definition both given	Copy hosted by Fraud Blocker, a click-fraud vendor, with a Fraud Blocker chapter; sponsor not stated	Forecast	◆ [84]
IAS, 20th Media Quality Report	$100 billion	2024	Cites Juniper's forecast	IAS, a verification vendor	Forecast cited as an estimate	◆
CHEQ with Roberto Cavazos (University of Baltimore)	$35 billion	2020	An assumed 10.5% "composite rate" times $332.84 billion of global digital ad spend (eMarketer)	CHEQ, a bot and click-protection vendor	Forecast	◆
ANA and White Ops (now HUMAN), Bot Baseline	$5.8 billion; $14 billion without anti-fraud measures	2019	Bot fraud only; 50 ANA members' campaigns, August and September 2018, projected globally	White Ops analysed the data	Projection; counterfactual	◆
TAG with the 4A's, ANA and IAB (Cunningham.Tech Consulting), US Ad Fraud Savings Report	$10.8 billion saved	2023	$11.78 billion counterfactual (9.96% unfiltered IVT from DoubleVerify, IAS, Oracle Moat and Pixalate times $118.2 billion of US display and video revenue) minus $979 million (0.92% times $106.4 billion through TAG-certified and non-certified channels)	The sponsors evaluate their own programmes	Counterfactual model	◆
TAG, European Ad Fraud Savings Report (release of 17 June 2025)	EUR 3.45 billion saved	2023	Counterfactual from a 10.36% unfiltered IVT rate; subtracts losses in protected and unprotected channels	Same developer; TAG evaluates its own programmes	Counterfactual model	[85]
DoubleVerify, release of 7 May 2026	About $1.8 million per billion CTV impressions	Not stated	Unprotected campaigns; price per thousand impressions (CPM), fraud rate and method not stated	DoubleVerify	Per-unit estimate	◆

They differ in definition, base and interest. Juniper's whitepaper gives a broad definition and a click-based one without saying which its forecast uses, CHEQ counts impressions, clicks, conversions and data events, and the ANA and White Ops counted bot fraud only ◆ ◆ ◆. IAS sizes the problem with Juniper's figure in its own report ◆. TAG's US saving subtracts actual losses only on the $106.4 billion that ran through certified and non-certified channels, so losses on the rest of the $118.2 billion, at the report's own unfiltered rate, count as saved; the European report, by the same developer, subtracts both ◆ [85].

Industry-wide ad fraud estimates, as published

Sponsored and commissioned estimates with their definitions and producers. Values are as published; none is the paper's own, and no two can be added or compared as if they measured one thing.

Machine-readable copy: data/estimates.csv.

PRODUCER	PERIOD	WHAT IS ESTIMATED	VALUE	YEAR	GEOGRAPHY	DEFINITION AND SCOPE	SPONSOR AND INTEREST	COMBINABLE?	SRC
Juniper Research	2023	ad fraud loss	$84.2 billion	2023	global	ad fraud defined as intentionally repeated clicking on PPC ads; 22% of digital ad spend	whitepaper copy hosted by Fraud Blocker with a Fraud Blocker chapter; sponsor not stated	No: each uses its own definition, scope and year	[84]
Juniper Research	2028	ad fraud loss	$172.3 billion	2028	global	same	same	No: each uses its own definition, scope and year	[84]
Juniper Research	2024	ad fraud loss	almost 100	end of 2024	global	same	same	No: each uses its own definition, scope and year	[84]
Juniper Research	2023	spend recoverable by fraud mitigation platforms	$23 billion	2023	global	same	same	No: each uses its own definition, scope and year	[84]
Juniper Research	2028	spend recoverable by fraud mitigation platforms	$47 billion	2028	global	same	same	No: each uses its own definition, scope and year	[84]
IAS, MQR 20th edition	2024	ad fraud cost cited by IAS	$100 billion	2024	global	cites Juniper (September 2023)	IAS (verification vendor)	No: each uses its own definition, scope and year	[187]
CHEQ with Prof. Roberto Cavazos	2020	direct cost of ad fraud	$35 billion	2020	global	impressions, clicks, conversions and data events	CHEQ (bot and click protection vendor)	No: each uses its own definition, scope and year	[188]
ANA and White Ops	2019	bot fraud losses	$5.8 billion	2019	global	bot fraud; 50 ANA members' campaigns, August to September 2018	White Ops (now HUMAN) analysed the data	No: each uses its own definition, scope and year	[189]
ANA and White Ops	2017	bot fraud losses (previous study)	$6.5 billion	2017	global	bot fraud	same	No: each uses its own definition, scope and year	[189]
ANA and White Ops	2019	bot fraud losses absent anti-fraud measures (counterfactual)	$14 billion	annual (2019 study)	global	bot fraud	same	No: each uses its own definition, scope and year	[189]
TAG with 4A's, ANA, IAB (Cunningham.Tech)	2023	IVT losses absent anti-fraud programmes (counterfactual)	$11.78 billion	2023	US	display and video; 9.96% x $118.2bn	the sponsors' own programmes	No: each uses its own definition, scope and year	[190]
same	2023	IVT losses in protected channels (modelled)	$0.979 billion	2023	US	0.92% x $106.4bn	same	No: each uses its own definition, scope and year	[190]
same	2023	savings attributed to anti-fraud programmes (counterfactual difference)	$10.8 billion	2023	US	11.78 - 0.979	same	No: each uses its own definition, scope and year	[190]
DoubleVerify	not stated	fraud cost per billion unprotected CTV impressions	$0.0018 billion	not stated	global	CTV; 'conservative estimates'; method not stated	DV (verification vendor)	No: each uses its own definition, scope and year	[191]

CHAPTER 4

Whether a verdict decides money

A third-party report determines money in 2 of 16 readable entries on the fixed platform list, or 12.5%, in the 10 November 2026 coding ◆. The cases are The Trade Desk's publisher terms and Microsoft Invest's route, whose end was announced for February 2026 ◆ ◆. No paid settlement on a third-party verdict or matched price comparison was found ◆ ◆. H2 is supported on that restricted reading, while coding all contracts in its registered population can cross the contradiction threshold ◆.

IN SHORT

Of 17 platform entries fixed at registration, the initial screen read 16 ◆. The Trade Desk makes accredited reports conclusive for publisher fees, and Microsoft Invest's refund route used an approved verifier's report ◆ ◆. These give the platform-only count of 2/16, or 12.5%, rising to 25% with review-qualified clauses ◆. Thirteen entries describe first-party validation ◆. Publicly recorded settlements and matched price differences were not found in the dated screen ◆ ◆. H2's broader registered population includes public contracts, whose inclusion and coding can move the share above 25% ◆. The buyer contract RM6123 gives an unnamed tool billing force between client and agency but expired in June 2026 ◆ ◆. The author retains H2 as supported under the registered listed-terms operational rule: 2/16 readable listed entries, or 12.5%. The broader population and treatment of review-qualified or trigger-only clauses remain reported sensitivities, not an unconditional claim across all public contracts ◆.

The claim under test

Two of the author's essays set out the framing under test. "Where the Surplus Settles" (September 2026) says "A DoubleVerify verdict moves nothing": verification is "a procurement record rather than a price input", and an impression clears at the same cost per thousand impressions (CPM) verified or not. "The Risk You Can Price" (August 2026) says advertising has auditors and accreditation bodies but no adjudicator with published rules of evidence and the power to bind an escrow, a claim about disputes over outcomes and conversions. Both are rival framings under test, never evidence. At registration one public contract already contradicted the first: The Trade Desk's OpenPath terms ◆.

H2's registered list covers 17 buying platforms and exchanges plus any public agency or advertiser contract found. It records a metric setting billable quantity, a third-party report deciding fees or remedies, an actual settlement on a verdict, and a matched price difference. First-party validation is recorded separately (Chapter 1). Third-party pre-bid targeting

changes eligible supply without itself setting a refund ◆. Both essay framings remain claims under test.

Sixteen of seventeen terms could be read

Sixteen of the seventeen entries could be read on 10 November 2026 ◆. Index Exchange's exchange-quality and policy pages carry no buyer or seller billing terms [86]. The verifier pass (a separate AI agent told to refute each finding the research agents made under the author's direction) later found Index's publisher-facing Marketplaces Terms of Use, last modified 13 May 2026, which pay publishers on Index's own records and give no third-party report a role [87]. The count keeps Index unreadable, the cautious choice, since reading those terms in would lower the share; one entry in seventeen is far from the third that would leave H2 unresolved.

Readable terms almost all bill on the platform's own count. Google Ads advertisers buying viewable CPM (vCPM) on Display Network campaigns pay for impressions that Google's own Active View counts as viewable (a small portion of served impressions might also be charged), and DV360 bills outstream video only on viewable impressions; DV360's viewable CPM is otherwise a bid goal, and its pay-per-viewable insertion orders were archived on 1 August 2023 ◆. Google Ads does not charge for invalid traffic (IVT) detected before billing and bases charges solely on Google's metrics ◆. DV360's exchanges refund traffic Google considers invalid, with HUMAN (formerly White Ops) as an extra check feeding Google's decision ◆. Amazon bases fees solely on its own measurements, and its traffic validation removes IVT detected within 30 days from invoices ◆. Meta does not charge for impressions it determines invalid ◆. Reddit's terms rule out charges based on third-party reporting, even from its approved measurement partners ◆, and Magnite's own determination is dispositive once a buyer's dispute stays unresolved for 30 days ◆. Yahoo's measurements are definitive ◆, and Snap, Pinterest, X and LinkedIn bill on their own ◆.

In 13 of the 16 readable entries the pages read describe the platform's own check of validity or viewability: first-party validation. For TikTok, Pinterest and X they describe none, though all three bill on their own data, and TikTok's terms make no commitment about invalid activity ◆ ◆. Third parties' delivery counts do carry billing force in places, in Google's discrepancy credit and TikTok's 20% review, but those count delivery, not verdicts ◆ ◆.

Where a third-party report decides

OpenPath clause 4.4 in v.10.31.24, unchanged from clause 4.5 in v.7.5.23, makes MRC-accredited reports used by The Trade Desk or its Ad Partners conclusive for fees owed to publishers [88] ◆. Invalid Inventory is unpaid or refunded within 30 days of notice, with a 1% monthly charge available for repeat sources ◆. The supplier terms carry the same mechanism ◆. Advertiser terms instead use The Trade Desk's own metrics and give no corresponding refund right ◆. The screened records show neither invocation frequency nor named verifier reports ◆.

Microsoft Invest's April 2025 documentation requires monthly claims backed by an approved partner's report, refunds the claimed quantity at Microsoft's average CPM, and adjusts seller accounts ◆. No partner or binding policy text was found ◆. Sellers retain a

dispute route, which could lower the strict numerator by one ◆. Microsoft announced the end of Invest buying for 28 February 2026, though pages still live on 10 November did not confirm retirement [89, 90].

PubMatic needs a vendor's report but keeps the decision. Section 8.4 of its Demand Partner Master Services Agreement (page modified 18 August 2026) credits impressions deemed fraudulent on an MRC-accredited vendor's scoring, within the window and threshold tabled below, only if PubMatic "verifies in good faith" that they were suspect; no make-good applies to viewability or brand safety ◆. The Trade Desk's SIVT Credit Program for advertisers worked the same way: claims needed a report from an MRC-accredited vendor such as Integral Ad Science (IAS) or DoubleVerify, the supplier's approval decided, and since late 2023, when its description was withdrawn, its status is unknown ◆.

Yahoo's display terms, in the same document as its DSP terms, also govern managed Yahoo DSP buys placed by insertion order ◆. They bill viewability-guaranteed vCPM line items on the viewable share measured by Yahoo or by an approved vendor the terms call "certified by the MRC", with Yahoo reserving a right to dispute vendor measurements and 10 days to reconcile; mobile and connected TV are among the exclusions. Some view-priced items may be invoiced on an approved vendor's count with no stated dispute right, which, folded in, would add one to the strict count ◆.

Coded strictly, a third-party report has determining effect in 2 of the 16 readable listed entries, 12.5%; set aside Microsoft's route, whose end Microsoft announced, and only The Trade Desk's clause remains ◆. Adding PubMatic's trigger-only credit and Yahoo's dispute-qualified vCPM clause, folded into the Yahoo DSP entry, takes the share to 25%, exactly a quarter ◆. The figure codes every entry, and the templates beside them, by component.

DOES A VERIFICATION VERDICT DECIDE WHAT IS PAID? THE REGISTERED LIST AND THE PUBLIC TEMPLATES

	Platform bills or credits on its own measurement	A third-party count sets the billable quantity (a)	A third-party report decides a fee or credit (b)	A settlement paid on a third-party verdict (c)	Price evidence, verified vs unverified (d)
Google Ads	yes	no	no	none found	none found
Display & Video 360	yes	no	no	own detection	none found
Google Ad Manager	yes	n/a	no	none found	none found
Amazon Ads	yes	no	no	none found	none found
Meta	yes	no	no	none found	none found
TikTok	not stated	no	no	none found	none found
Snap	yes	no	no	none found	none found
Pinterest	not stated	no	no	none found	none found
Reddit	yes	no	no	none found	none found
X	not stated	no	no	none found	none found
LinkedIn	yes	no	no	none found	none found
The Trade Desk	yes	optional	yes	none found	none found
Microsoft Invest	yes	no	yes (retired)	none found	none found
Yahoo DSP	yes	no	no	none found	none found
Magnite	yes	no	no	none found	none found
PubMatic	yes	no	trigger only	none found	none found
Index Exchange	n/e	n/e	n/e	n/e	n/e
◦ ANA media-buying template v3.0	n/a	optional	yes	n/a	n/a
◦ IAB Terms, Direct Buy Addendum	n/a	optional	no	n/a	n/a
◦ IAB Measurement Addendum	n/a	optional	no	n/a	n/a
◦ 4A's/IAB Standard Terms v3.0	n/a	optional	no	n/a	n/a
◦ Disney DAS IAB Addendum	n/a	no	qualified	n/a	n/a
◦ Yahoo Display Terms	yes	optional	qualified	n/a	n/a
◦ Spotify Master Terms	yes	no	no	n/a	n/a

Yes, in the terms read · Yes, in a route since retired · Optional or by agreement
Starts a review; the platform decides · Decides subject to a dispute right · No
None found · Settlement on the platform's own detectic · Not stated
Terms not found · Not applicable (template)

◦ marks public templates and contracts outside the list fixed at registration. The Trade Desk's clause binds it and publishers (supply side), not advertisers; Microsoft retired Invest buying on 28 February 2026. A trigger-only clause starts a review that the platform decides. Read between 9 and 10 November 2026; Index Exchange published no billing terms that were found.

A third-party report decides money in two of sixteen readable terms, and no settlement was found. The platform and exchange terms on the list fixed at registration (hypothesis H2), with the public templates read alongside, coded for the four components: whether a third-party count sets the billable quantity, whether a third-party report decides a fee or credit, whether a settlement paid on a third-party verdict is documented, and whether verified and unverified impressions of the same inventory are shown to clear at different prices. Every large platform bills or credits on its own measurement.

Units: coded state per term · Scope: 17 listed terms (16 readable) and 7 public templates and contracts, read 9 to 10 November 2026 · Sources [192] [193] [194] [195] [196] [197] [198] [199] [200] [201] [202] [203] [204] [205] [206] [207] [208] [209] [210] [211] [212] [213] [214] [215] [216] [217] [88] [218] [219] [220] [89] [221] [222] [223] [224] [225] [226] [227] [228] [229] [230] [231] [86] [232] [233] [234] [235] [236] [237] [238] [239] [240] [241]

The share moves with coding: 12.5% on the strict count, up to 25% if trigger-only and qualified clauses are counted. Executed contracts are private; this is what public terms say.

Windows, thresholds and who decides

The terms also set claim windows and who rules. On the list, the platform rules everywhere except where a report is conclusive (The Trade Desk's supply-side clauses) or sets the refunded quantity (Microsoft Invest) ◆ ◆. The windows start from different points and are not one series.

TERM	MECHANISM	WINDOW	THRESHOLD OR FORMULA	WHO DECIDES	EVIDENCE
Google Ads	Advertiser-requested IVT investigation	Activity in the last 60 days	None stated	Google	◆
Google Ads (US terms of 29 July 2025 and 1 July 2026)	Third-party ad server discrepancy credit	Claim within 60 days of the invoice	Gap above 10%; credit of (90% of Google's count minus the ad server's) times the average CPM or vCPM	Google	◆
Google Ad Manager	Advertiser credits behind publisher deductions	Traffic in the 60 days before the invoice	None stated	Google; no appeal	◆
Amazon Ads	IVT monitoring; fee claims	Up to 30 days; claims within 60 days of the charge	Gap of up to 5% from third-party IVT rates expected	Amazon	◆
TikTok	Third-party count review (reservation display ads); payment disputes	Disputes waived a year after the charge	Review only above a 20% difference; TikTok's data final if no error is found	TikTok, in a joint review	◆
Pinterest; X	Charge disputes	60 days	None stated	The platform	◆
LinkedIn	Fee disputes and invalid-activity credit claims	90 days	Credit at LinkedIn's sole discretion	LinkedIn	◆
Reddit	Ads Credits claim period	30 days	Credits at Reddit's discretion	Reddit	◆
Yahoo	Charge disputes; vCPM vendor reconciliation	30 days of the invoice; 10 days	Viewability minimum set per order	Yahoo; the parties for vCPM	◆ ◆
Magnite	Notice of invalid activity	10 days after the invoice; Magnite decides after 30 more	Reasonable evidence	Magnite	◆
PubMatic	Noncompliant Impression report	30 days after month-end	Over US$500 a month; capped at winning bids	PubMatic	◆
The Trade Desk (OpenPath, supplier terms)	Seller refund of paid Invalid Inventory	30 days after The Trade Desk's notice	1% of the month's payable for repeat sources	The reports, deemed conclusive	◆ ◆
The Trade Desk (advertiser terms)	Supply discrepancy	30 days of the invoice	The Trade Desk bears gaps above 10%	The Trade Desk	◆
The Trade Desk (SIVT Credit Program, withdrawn 2023)	Client SIVT claim	30 days of the incident	Over US$1,000 of media spend	Supplier approval	◆
Microsoft Invest (end announced for 28 February 2026)	IVT refund claim	10 days after month close	None stated	The partner's report sets the quantity	◆

Templates and contracts outside the list

The population also takes in public contracts: templates from the ANA (Association of National Advertisers), the IAB (Interactive Advertising Bureau) and the 4A's (American Association of Advertising Agencies), sellers' own terms and one government framework.

The paper reads them beside the list; the verifier pass argued they belong in the test, and the verdict below reports both readings ◆.

DOCUMENT	DATE	WHAT IT DOES WITH A VERIFICATION VERDICT	CODED (B)	EVIDENCE
ANA Master Media Buying Services Agreement Template v3.0	June 2023	The agency must make media owners accept the advertiser's ad-verification tool as the basis of the billable metric, and seek terms under which the advertiser does not pay for IVT, brand-unsafe or non-viewable placements the tool identifies	Yes, in template language	◆
IAB General Terms v1.1, Direct Buy Addendum v1.0, Appendix 1 v1.1	Released 18 May 2026	Controlling measurement set per order; an ad verification company is mutually agreed; remedies mutually chosen	No default	◆
IAB Measurement Services Addendum v1.0	Posted as final 28 October 2026; the file reads "DRAFT"	Provides for a provider's outputs designated as the controlling measurement under another agreement; sets no default	No default	◆
4A's/IAB Standard Terms v3.0	2010	A guideline-compliant third-party ad server's delivery count controls invoices; adjacency remedy on written request, with no named decider; no verification term	No	◆ ◆
OpenX addendum to the 4A's/IAB terms	Undated	On an unresolved discrepancy, pays on the average of the two measurements; no verification term	No	[91]
Disney Advertising addendum, US media buys	28 March 2023	Verifier reports decide adjacency compliance; Disney reviews the incident list before any remedy; its own ad server controls billing	Qualified	◆
Yahoo Display Advertising Program Terms, clause 10.1	1 January 2026	vCPM billed on Yahoo's or an approved vendor's viewable share, subject to Yahoo's dispute right	Qualified	◆
Spotify Master Advertising Terms	1 July 2026	Spotify's metrics control over third-party statistics unless an order says otherwise	No	◆
RM6123 Media Services framework (UK public sector)	14 December 2021 to 13 June 2026 (expired)	Where ad fraud tools are used, the client pays only for what the tool verifies as human traffic	Yes, between client and agency	◆

The ANA template is the strongest buyer ask in writing, but no member has to use it, its non-payment clauses bind the agency only to commercially reasonable efforts, and whether media owners accept them is not public ◆.

No settlement and no price difference found

In the sources screened by 10 November 2026, no settlement paid on a third-party verdict was found ◆. The one advertiser settlement over invalid-traffic credits found, in AdTrader v. Google, was given final approval on 1 November 2022: a non-reversionary US$7 million fund for US AdWords (now Google Ads) advertisers who were not refunded for ads on Google AdX publishers' sites although Google withheld those publishers' payments for invalid activity or other breach. The money moved on Google's own determinations, and a settlement is not a finding of liability ◆. Uber's US$6 million settlement with Phunware (9 October 2020), over fraud claims, names no verification vendor ◆. The screen was thin, with no docket search by vendor name, and credits paid privately would not appear in public records ◆.

Parties say verdicts move money; none documents a payment. DoubleVerify's May 2025 complaint against Adalytics, a party's pleading, says general invalid traffic (GIVT) is deducted from the billable count before a customer is billed and SIVT is flagged so the advertiser can seek compensation from the seller; it cites no contract, platform term or

settlement, and MRC says it has no insight into billing arrangements ◆. A 2019 release from the ANA and White Ops (now HUMAN) said most fraudulent impressions were by then invalidated by platforms, filtered before payment or recovered through claw backs, but it recorded no claw back ◆. Pixalate has sold a signed Clawback Report for refund negotiations since 4 November 2025, but its announcements name no counterparty or amount refunded ◆. The verifiers' own terms promise no more: IAS's and DoubleVerify's make their statistics determinative, or potentially so, for billing their own fees, and Pixalate's bar clients from using its output in litigation or advertising claims without its consent ◆ ◆ ◆.

Component (d), the essay's "same CPM", is neither shown nor refuted: no public evidence found by 10 November 2026 compares the prices of verified and unverified impressions of the same inventory ◆. The closest are publishers' statements to Digiday in December 2017, one reporting only a 13% CPM gap between its most and least viewable units, another, anonymous, roughly estimating CPMs about 20% higher per 30-point gain in viewability; those compare viewability levels, not verification status, and billing on viewable impressions makes a verdict the unit paid for, not a different price for the same impression ◆. The ANA's TrueCPM, $7.46 for higher-performing and $19.04 for lower-performing advertisers per thousand quality-tested impressions (27 May 2026), compares advertisers, not inventory [69, 92].

The verdict on H2

H2 requires third-party determining effect in no more than 25% of readable listed terms, with no documented paid settlement or matched price difference. Its registered population also includes public agency or advertiser contracts. The platform-only reading meets the support test at 2/16, or 12.5%, and 25% on lenient coding ◆. Across all contracts found, strict coding gives 4/26 to 4/23, while including review-qualified or trigger-only clauses can exceed 25% ◆. The rule counts the listed terms, so H2 is supported on its registered test; the wider contract population and the qualified clauses are reported as sensitivity, not as the test ◆.

Three qualifications travel with it. The rule names no tier for clauses that keep a review or dispute right, a gap the verifier pass flagged along with the treatment of ended terms; on the list the verdict holds at every tier, and dropping Microsoft's route only lowers the share ◆. Across every public contract found as well, RM6123, the OpenX addendum and Index's publisher terms among them, the share stays under a quarter on strict coding, falls either side of it when review-qualified clauses are counted, and passes it when verification-gated triggers are added; on terms in force at reading it stays under a quarter at every coding ◆. And the party statements above document no payment: DoubleVerify's pleading describes billing mechanics, which H2 counts as neither (b) nor (c), and the 2019 release asserts recoveries without a verdict source, payer or amount; were it read as evidence of (c), the rule would leave (c) unresolved ◆ ◆.

What buyers write down

Buyers document verification mostly as demands, guidance and certification duties ◆. The one public buyer contract found that gives a tool's verdict determining effect is the UK

government's RM6123 Media Services framework (14 December 2021 to 13 June 2026). Where ad fraud tools are employed, the public-sector client pays only for impressions, clicks or conversions the tool verifies as human traffic, and the agency must use reasonable endeavours to recover the value of fraudulent, non-viewable or unsafe ads it paid for ◆. The clause binds client and agency, not the seller; it names no vendor, and the tool need not be independent or accredited ◆. No record shows money withheld or recovered under it, and "human" is the contract's word: an IVT tool classifies an impression it does not filter as valid ◆. On 10 November 2026 the government's upcoming-agreements page listed no media-buying successor [93].

The rest is demand. The Global Media Charter of the World Federation of Advertisers (WFA), dated 21 June 2018, asked for refunds of all media investment associated with invalid traffic, fees and commissions included ◆. The ANA's Complete Report (December 2023) tells advertisers not to pay for invalid, unmeasurable or non-viewable impressions and to settle refund processes with DSPs in advance ◆. ISBA, the UK advertisers' body, tells advertisers to ask whether they pay for blocked impressions ◆. Member-only templates could not be read, so their clauses are unknown, not absent ◆.

The ANA's study used vendors' classifications for measurability, viewability and IVT (Chapter 3) ◆. Its Q3 2025 release called the new risk measure an AI-powered standard without naming the classifier, while the benchmark report identified Mobian and reported 99.1% of spend in low-risk environments ◆ ◆. Those are measured outcomes without billing force.

Who pays, and what the seals certify

Advertisers pay DoubleVerify per thousand impressions, often through agencies or platforms ◆. IAS pays DSPs a share of revenue earned through their platforms ◆. The ANA's 2023 member survey reports 78% holding direct verification contracts but publishes no response count ◆. No verification rate card was found in that screen ◆. Chapter 7 covers seller revenue and interests.

Every holding company in scope shows TAG seals for a buying entity, but a seal certifies process, not a verifier. TAG's Brand Safety Certified Guidelines v5.0 (July 2026) require certified buyers, agencies included, to define make-good and clawback procedures and to screen all transactions with content verification services that have passed independent validation or with inclusion and exclusion lists, so lists alone can meet the rule ◆. TAG's Certified Against Fraud rules add a duty to filter all transactions for invalid traffic ◆. The registry read on 10 November 2026 shows:

HOLDING COMPANY	REGISTRY ENTITY	BRAND SAFETY CERTIFIED	CERTIFIED AGAINST FRAUD	EVIDENCE
WPP	WPP Media	Independent validation; Asia Pacific, UK, US, Ireland	Self-attested; North America	◆ [94]
Publicis	Publicis Groupe/MMS USA Holdings	Independent validation; global	Independent validation; global	[94]
Omnicom	Omnicom Media Group	Independent validation; UK	Self-attested; North America, Asia Pacific, Africa, Central and South America, Middle East, Oceania, Russia	◆ [94]

HOLDING COMPANY	REGISTRY ENTITY	BRAND SAFETY CERTIFIED	CERTIFIED AGAINST FRAUD	EVIDENCE
Omnicom (IPG since 26 November 2025)	IPG Mediabrands	Independent validation; global	Independent validation; global	◆ [94]
dentsu	Dentsu Aegis Network	Independent validation; Ireland, UK	Independent validation; global	[94]
Havas	Havas Media	Independent validation; UK, Ireland	Independent validation; UK	[94]

The registry's names lag corporate change (IPG Mediabrands, Dentsu Aegis Network) ◆. Omnicom's and WPP's FY2025 annual reports say nothing about brand safety, viewability, invalid traffic or ad fraud ◆.

Disputes on the record

Both public verifiers' filings acknowledge billing disputes without naming any, and on 10 November 2026 CourtListener's RECAP index showed no suit by an advertiser or agency against either by case name, a search that misses sealed matters, arbitrations and state courts ◆. The register's other items record no payment on a third-party verdict (Chapter 9):

DATE	ITEM	WHAT IT IS	POSTURE	EVIDENCE
3 June 2020	Sizmek DSP v. DoubleVerify (Bankr. S.D.N.Y., Adv. No. 20-01119)	Adversary complaint by a buy-side platform in bankruptcy	Terminated 1 March 2021; no document public	◆
3 June 2020	Sizmek DSP v. Integral Ad Science (Adv. No. 20-01134)	Adversary complaint by a buy-side platform in bankruptcy	Closed by so-ordered stipulation on 12 April 2021; no document public	◆ [95]
6 August 2024	Rumble v. WFA, Diageo, WPP plc and GroupM Worldwide (N.D. Tex.)	Antitrust complaint by a publisher	Dismissed without prejudice on 13 August 2025 for lack of personal jurisdiction and improper venue, a procedural ruling	◆
7 August 2024	Adalytics report on brand-safety technology	Independent report relaying statements	Relays that a brand was told its campaigns were 100% brand safe; records no contract, refund or dispute outcome	◆
7 February 2025	Adalytics report on an image host	Independent report relaying statements	An intention to seek refunds stated; none documented	◆
26 September 2025	Federal Trade Commission Decision and Order, Omnicom and IPG (Docket C-4823)	Consent order, a settlement and not a finding	Bars the merged agency, in the US, from steering spend or relying on exclusion lists on "Covered Bases" such as political viewpoint; client-directed lists allowed; compliance monitor	◆

What the finding means for the framing

The public terms do not bear out "moves nothing" as a general claim: a third-party report decides fees in The Trade Desk's supply-side terms, decided refunds in Microsoft Invest's route, whose end was announced, sets vCPM billing in Yahoo's display terms subject to dispute, and is the billing basis in the ANA's template and the UK's expired RM6123 framework ◆ ◆ ◆ ◆ ◆. As a description of the largest platforms' own billing terms, the line holds: Google, Amazon, Meta, TikTok, Snap, Pinterest, Reddit, X and LinkedIn bill on their own counts, and any verdict that moves money there is their own ◆ ◆ ◆. The "same CPM" half is neither shown nor refuted ◆. On who decides, the terms read name the platform, the parties or a report deemed conclusive, and none names an outside body that hears disputes

under published rules of evidence; that fits the adjudicator essay as far as these terms go, though the essay concerns conversion disputes ◆ ◆. All of this rests on public terms read on 10 November 2026, to be rechecked at the freeze; executed contracts, negotiated orders and privately paid credits were outside reach.

CHAPTER 5

Inside the platforms

Where the nine platforms document a route, third-party verification mostly uses platform-supplied data ◆. Meta supplies filtered impressions, YouTube routes partners through Ads Data Hub, and Amazon runs vendor detection inside its environment ◆ ◆ ◆. TikTok's documented OM SDK route contradicts H4 under the registered comparator ◆. That conclusion carries medium confidence because the route was described in 2021–2022 and TikTok controls the integration ◆.

IN SHORT

Meta, YouTube and Amazon document platform-supplied routes, with different data limits ◆ ◆ ◆. Snap and LinkedIn document trackers without establishing where they fire, while X, Pinterest and Reddit leave the route unstated ◆. Four of nine platforms have accredited third-party measurement, one has third-party content-level measurement and six have accredited first-party measurement, on MRC's 10 November 2026 list ◆. Meta and YouTube third-party viewability is accredited at GIVT level ◆ ◆. Documented pre-bid controls are executed by the platform ◆. Post-bid blocking remains undocumented outside Amazon's unresolved tag scope ◆. H4 is contradicted on TikTok under the literal tag comparator, at medium confidence ◆ ◆. Platform-reported pass rates remain bounded by the impressions and content made measurable ◆.

How a route is coded

H4, registered on 9 November 2026, predicts that on each of the nine platforms (Meta, YouTube, TikTok, Amazon, X, Snap, Pinterest, Reddit and LinkedIn), documented third-party brand-safety, suitability and viewability measurement of the platform's own placements runs on data the platform supplies, "not on the verifier's own code or crawler observing the placement" ◆. Its comparator is "the verifier's own tag or SDK running in the placement, or its own crawl of the content adjacent to the ad" ◆.

Data access limits what partners can check. MRC says IAS lacks user- and device-level data on Meta and cannot independently filter IVT there ◆. DoubleVerify's FY2025 filing says integration partners control its platform solutions and have supplied inaccurate signals ◆. It says it uses its own tags when possible, while IAS says methods vary by platform ◆ ◆.

AI research agents working under the author's direction read the platforms' help, policy and developer pages, the verifiers' integration pages and MRC's records on 9 and 10 November 2026, and a separate AI verifier, told to refute each finding, re-checked it against its sources in what this paper calls the verifier pass (Chapter 13). Each platform, metric and verifier was coded with one of four routes: platform-supplied (a feed, interface, log, clean room or

server-to-server processing inside the platform's environment), verifier-executed (the verifier's own tag, SDK or crawler), a vendor tracker the platform accepts, or not stated, which is the code for a partner named without a route ◆. Off-platform supply that a platform sells, such as LinkedIn's Audience Network, is outside H4 ◆. Every entry in the table is bounded to the pages screened.

PLATFORM	PARTNERS DOCUMENTED	VIEWABILITY AND INVALID TRAFFIC	BRAND SAFETY AND SUITABILITY	EVENT-LEVEL OR AGGREGATE	EVIDENCE
Meta	DoubleVerify, IAS, Zefr and Adloox, on Meta's pages	Platform-supplied: impressions Meta measures, with Meta's invalid-traffic removal	Platform-supplied on Feed and in-stream video, inferred from MRC's description of Meta's pipeline (under review); not stated for Reels proper, Stories or Threads	Impression-level without user or device identifiers; per-content impression estimates, hourly	◆ ◆ ◆ ◆ ◆
YouTube	DoubleVerify, IAS and Zefr, on Google's pages	Platform-supplied through Ads Data Hub; no third-party pixels accepted	Exposure through Ads Data Hub; how vendors obtain the content not stated	Privacy-checked aggregates of roughly 10 to 50 users a row (inference)	◆ ◆ ◆ ◆
TikTok	DoubleVerify, IAS and Zefr, on TikTok's pages	Verifier-executed: the vendor's tag through TikTok's OM SDK (2021, 2022); not restated in 2026	IAS: the same integration (2022), content route not stated; DoubleVerify and Zefr: not stated	Not stated	◆ ◆ ◆
Amazon	DoubleVerify, IAS and Adloox, in Amazon's policy	Platform-supplied: server to server, detection inside Amazon's environment; monitoring tags also listed for supply not named	Server to server per Amazon's policy; scope not stated	Not stated	◆ ◆
X	IAS and DoubleVerify, on X's pages	Not stated	Not stated (DoubleVerify says it classifies the posts above and below ads)	Not stated	◆
Snap	IAS and DoubleVerify, and Zefr for brand safety, on Snap's pages	Approved vendors' 1×1 trackers, fired when the ad renders; where they fire from not stated	IAS trackers required; content route not stated	Not stated	◆ ◆
Pinterest	IAS and DoubleVerify, on Pinterest's and the vendors' pages	Not stated	Not stated	Not stated	◆
Reddit	IAS on Reddit's pages; DoubleVerify in its own release	Not stated	Not stated	Not stated	◆
LinkedIn	DoubleVerify, and IAS for video ads only, on LinkedIn's pages	Image or tracking tags; JavaScript and iFrame tags refused	No third-party measurement of the feed documented	Not stated	◆

The figure shows the coded routes beside MRC's status for third-party and first-party measurement on each platform ◆ ◆.

INSIDE THE NINE LARGE PLATFORMS: WHERE THIRD-PARTY MEASUREMENT GETS ITS DATA

	Viewability and IVT	Brand safety and suitability	MRC: third-party measurement	MRC: the platform's own measurement
Meta	platform data	platform data	accredited	accredited
YouTube	platform data	platform data	accredited	accredited
TikTok	verifier tag	verifier tag	accredited	
Amazon	platform data	platform data	accredited	accredited
X	not stated	not stated		
Snap	platform-fired tag	not stated		accredited
Pinterest	not stated	not stated		accredited
Reddit	not stated	not stated		
LinkedIn	platform-fired tag	not stated		accredited

Runs on data the platform supplies · Vendor pixel the platform fires · Verifier's own code in the placement
Route not documented · MRC-accredited · No MRC accreditation

TikTok's route rests on DoubleVerify's (2021) and IAS's (2022) documents describing measurement through TikTok's Open Measurement SDK; no 2026 page restates it, and TikTok controls the SDK (medium confidence). Snap's and LinkedIn's first-party accreditations cover the platforms' own measurement, not a verifier's. Read 10 November 2026.

> **On eight of nine platforms, third-party verification checks the platform's own data or cannot say how it gets it.** For each of the nine large platforms, the documented route by which third-party verifiers measure viewability, invalid traffic and brand safety on the platform's own placements, and whether MRC accredits third-party measurement there or only the platform's own. Only on TikTok is the verifier's own code documented as running in the placement.
>
> Units: coded route per platform and metric · Scope: Meta, YouTube, TikTok, Amazon, X, Snap, Pinterest, Reddit, LinkedIn; documents read 10 November 2026 · Sources [101] [67] [65] [242] [243] [244] [245] [246] [152] [247] [248] [249] [250] [251] [204] [66] [252] [253] [98] [97] [254] [255] [256] [257] [258] [99] [259] [260] [261] [262] [263] [264] [265] [266] [267] [268] [269] [270]
>
> A partner's name alone is coded "not stated". Platform pages change often; the routes are status-sensitive.

Meta, YouTube and Amazon supply the data

Meta's partners measure what Meta passes them. MRC accredits DoubleVerify's and IAS's reporting of Meta impressions and viewability only on placements that Meta itself measures and has accredited, net of GIVT, and IAS's letter of 9 July 2026 records that, because of limits Meta imposes, IAS cannot independently identify and filter invalid traffic there and has no user- or device-level data ◆ ◆. The French competition authority's interim-measures decision of 4 May 2023, an account of the parties' statements rather than a ruling on the merits, describes Meta supplying the data through a programming interface, removing the invalid traffic it identifies before passing impression-level data to viewability partners, and not letting partners collect raw data with their own technology ◆ ◆.

For suitability, Meta's announcement page dates its partners from Zefr (30 March 2023) to Adloox (24 October 2024) without saying what data it supplies ◆. MRC's under-review list, read on 10 November 2026, carries a Meta Third-Party Brand Suitability Pipeline, audit reviewed in August 2026: defined data fields that Meta delivers to partners for the text, image, audio and video in Facebook in-stream video and the Facebook and Instagram feeds ◆. Meta's Content Delivery Reports interface gives partners, for Facebook in-stream and Reels overlay ads, each public content item's ID with estimated impressions, hourly, over a 15-day lookback, and the IDs can be used to retrieve the video file; neither page names a verifier ◆ ◆. How partners obtain content on Reels proper, Stories and Threads is not documented in the pages opened ◆.

Google accepts no third-party tracking pixels for YouTube measurement, and enabling a vendor in Google Ads means choosing to share campaign data with it ◆. DoubleVerify's, IAS's and Zefr's accredited YouTube viewability is processing of Google's Ads Data Hub (ADH) data at GIVT level, while Google's own ADH accreditation covers its metrics including SIVT and the passthrough reporting that vendors generate by querying ADH, which MRC's report for the first quarter of 2026 calls templated vendor queries ◆ [62]. Google says ADH's privacy checks need roughly 10 to 50 unique users per result row and are designed to help prevent the transmission of data about individual users, so verifiers query event data Google holds without receiving event-level records; that is an inference, though Google's page names measurement partners among ADH's users ◆.

For content, exposure reaches Zefr through ADH, but how the vendors obtain the video they classify is not documented on the pages opened; MRC's chart calls Zefr's accredited method independent URL classification ◆. YouTube's derived-metrics policy, last updated on 14 September 2026, reserves the impression-level data needed for safety claims about specific impressions to Brand Safety Reporting Partners, whom it does not name ◆. Channel Factory says it uses YouTube's interface and its own engine to scan video, audio and metadata and build suitability lists, which this paper reads as a pre-bid list product ◆. Adalytics argued in a report public by 27 June 2023 that YouTube's measurement vendors rely on aggregated server-to-server transfers (Chapter 3) ◆.

Amazon says its server-to-server approach lets DoubleVerify and IAS apply their detection methods inside its secure environment, and its policy, last updated on 1 June 2026, offers DoubleVerify, IAS and Adloox on its demand-side platform (DSP) through server-to-server integrations and refuses JavaScript tags for third-party tracking in its mobile app ◆. The same policy lists JavaScript monitoring tags on creatives hosted by Amazon DSP, and monitoring-and-blocking tags on creatives from third-party ad servers, without naming the supply, so verifier code on some Amazon placements is not excluded ◆ ◆. Which data fields Amazon passes is not documented on the pages opened ◆. Amazon says its own systems, using signals from HUMAN (White Ops until March 2021), exclude identified invalid traffic from reports and invoices: a first-party verdict (Chapter 4) ◆ ◆.

TikTok runs the verifier's tag

TikTok is where the record differs. DoubleVerify (28 September 2021, read in a syndicated Business Wire copy) and IAS (9 May 2022) announced that they measure viewability and invalid traffic on TikTok's placements through TikTok's integration of IAB Tech Lab's OM SDK, and IAS said the same on 17 October 2022 of post-bid brand safety and suitability, in

markets it did not name ◆. In the OM SDK the app sends measurement signals and the provider's own tag collects them, which for viewability and invalid traffic is verifier-executed observation under H4's comparator ◆.

No 2026 source opened restates the OM SDK route, and TikTok controls which scripts load and what signals they receive ◆. Its current IAS-tag requirement supports continuity, but MRC separates DoubleVerify's TikTok traffic from its tag-based reporting ◆. The impression tag also leaves IAS's content-acquisition route unstated ◆.

TikTok's help pages offer DoubleVerify and IAS invalid-traffic and viewability measurement on the For You Feed and TikTok Ad Network (formerly Pangle); DoubleVerify's metrics there include Total Calls and Prerender Calls, names consistent with code calling home from the device but no proof of it ◆. MRC granted DoubleVerify initial accreditation for TikTok in-app video impressions and viewability, filtered for GIVT and SIVT, on 13 April 2026, and in December 2025 removed TikTok itself from in-process status because it did not intend to take part in a third-party audit ◆. None of TikTok's pages opened says what content data it gives its brand-safety partners ◆.

Snap and LinkedIn accept trackers

Snap documents approved tracking URLs, including a 1×1 impression tracker and IAS tags for suitability ◆. LinkedIn accepts image and tracking tags but rejects JavaScript and iFrames, with DoubleVerify's display 1×1 URL and an undescribed video tag ◆. Its suitability offer applies only to the off-platform Audience Network ◆.

Neither platform establishes whether the request fires from the device or its servers ◆. The verifier pass therefore treated the viewability route as unresolved, rather than accepting the research's platform-mediated inference ◆.

X, Pinterest and Reddit name partners only

X (formerly Twitter) names IAS and DoubleVerify on help pages read in a browser on 10 November 2026, for viewability and brand-safety measurement and for post-bid adjacency measurement in the Home Timeline, defining adjacency as one organic post above the ad and one below ◆. DoubleVerify said in January 2023 that it classifies those posts, but no page opened says how they reach the verifiers; that X must supply them is inference ◆. Pinterest, DoubleVerify and IAS document viewability and invalid-traffic measurement since 2020 and brand-safety reporting since 2024, with no route on any page opened ◆.

Reddit offers a Limited inventory tier, powered by IAS, to managed accounts, and DoubleVerify and IAS describe measurement on its feed and conversation placements ◆. Reddit's tracker page names ten approved partners, as some of those it works with, and neither verifier ◆. A trade report called DoubleVerify's integration server to server; no primary page opened confirms it ◆. Zefr announced a Reddit partnership on 30 September 2026 that no Reddit page opened names ◆.

What MRC accredits on the platforms

Accreditation on the platforms is MRC's, set letter by letter; it certifies a scoped measurement process against MRC's standards; the public accreditation letter does not supply the sold-category error-rate estimate defined in H3 (Chapter 2). On the list read on 10 November 2026, third-party measurement of a platform's own placements is accredited on four of the nine platforms (Meta, YouTube, TikTok and Amazon), content-level measurement only for Zefr on YouTube, and the platforms' own measurement on six ◆. A membership or a pre-assessment is not accreditation: X says it is an MRC member that completed a pre-assessment in December 2022 without a full audit, and MRC's chart of 22 October 2026 shows X as not in MRC audit ◆. Reddit appears on neither MRC list nor the chart ◆ ◆. Each row's filtration level, GIVT or SIVT, is the invalid-traffic standard its letter applies ◆ ◆ ◆.

PLATFORM	THIRD-PARTY MEASUREMENT ACCREDITED	THE PLATFORM'S OWN ACCREDITED MEASUREMENT	EVIDENCE
Meta	DoubleVerify: display and video impressions and viewability, desktop, mobile web and in-app, net of GIVT, Facebook and Instagram feeds only (13 April 2026). IAS: the same metrics net of GIVT on Facebook Feed in all three environments and, in-app, on Instagram Feed and Stories and Facebook in-stream video; Reels and Audience Network excluded (9 July 2026)	Facebook and Instagram impressions and video plays on desktop, mobile web and in-app, and first-party viewability report metrics (in-app video on MRC's list), by placement, including SIVT filtration (2 October 2026); content-level controls revoked in September 2025 on Meta's withdrawal	◆ ◆ ◆
YouTube	DoubleVerify (13 April 2026), IAS (9 July 2026) and Zefr (6 November 2025): Ads Data Hub video impressions and viewability, desktop, mobile web and in-app, GIVT level (IAS's higher level expressly not SIVT), in-stream and bumper; Shorts, in-feed and masthead excluded. Zefr: content-level labelling and reporting of English-language in-stream video and music video, desktop, mobile web, in-app and connected TV (CTV); Shorts, YouTube TV, Google TV, masthead, Google Video Partner videos and volumetric data excluded (20 January 2026)	Ads Data Hub metrics for in-stream YouTube and Google Video Partner inventory, desktop, mobile web and in-app, including SIVT, with third-party passthrough reporting (20 January 2026); Inventory Modes content-level controls on in-stream video and Shorts sold through Google Ads and Display & Video 360, desktop, mobile web and in-app; masthead, YouTube Kids and live streams excluded (28 May 2026)	◆ ◆ ◆
TikTok	DoubleVerify: video impressions and viewability, in-app, filtered for GIVT and SIVT (initial grant, 13 April 2026)	None; removed from in-process status in December 2025	◆
Amazon	IAS, server to server: display and video impressions, viewability, tracked ads and invalid traffic, desktop, mobile web and in-app, for Amazon Retail, IMDb and third-party supply; Fire TV and Twitch excluded; SIVT in the letter's body, net of GIVT in its enclosure (7 November 2025)	Amazon DSP display impressions and click-throughs for Amazon Retail, IMDb, Amazon Publisher Services and third-party exchanges, net of GIVT, and a first viewability grant across those sources and Alexa devices, net of GIVT (29 May 2026); Twitch and Prime Video video ads neither accredited nor in audit	◆ ◆
X	None; platform not in MRC audit	None; pre-assessment in December 2022 and no full audit, by X's account	◆
Snap	None accredited or in process	Impressions, viewability, swipe-ups, reach and video duration, in-app, GIVT (29 July 2026); desktop web denied in September 2025	◆
Pinterest	None accredited or in process	Display impressions and pin clicks (outbound clicks excluded), desktop, mobile web and in-app, GIVT (20 January 2026)	◆
Reddit	No entry	No entry	◆
LinkedIn	None	Impressions and clicks on LinkedIn.com and its app, desktop, mobile web and in-app, net of GIVT; Audience Network excluded (28 May 2026)	◆ [96]

On Meta and YouTube, where the verifiers process data the platform supplies, their accredited viewability stops at GIVT while the platforms' own measurement includes SIVT, and MRC's IAS letter gives the reason for Meta ◆ ◆ ◆. The two third-party accreditations on the nine platforms that include SIVT are DoubleVerify on TikTok, where the route documented in 2021 is the verifier's tag in the app, and IAS on Amazon, where detection runs inside Amazon's environment ◆ ◆. A withdrawal such as Meta's is not a finding on accuracy ◆.

MRC's own records also disagree. Its chart of 22 October 2026 footnotes Amazon DSP third-party viewability as accredited for DoubleVerify and in process for IAS, the reverse of the list and letters read on 10 November 2026, which show IAS accredited and no DoubleVerify Amazon entry ◆. The chart also shows DoubleVerify accredited on Instagram Stories and Facebook in-stream Reels video, beyond its feeds-only letter; MRC says the chart does not reflect each platform's full scope, so the letters and list are treated as the operative record ◆.

MRC's 2018 content-level supplement makes platform cooperation a condition of third-party content-level accreditation: the platform allows direct third-party content-level measurement or provides validated first-party data ◆. Two platform data services were under MRC review on 10 November 2026: Meta's suitability pipeline and Google's YouTube Relay Log, whose aggregated views and watch time for measurement partners are audience data rather than verification ◆ ◆. IAS's YouTube brand-safety classification has been in audit since April 2023 ◆.

Lists before the bid, little after it

Every third-party pre-bid control whose mechanism is documented is a verifier-supplied list, classification or tier that the platform executes, such as TikTok's partner-managed video exclusion lists, which take up to 200,000 videos at a time, and Meta's partner block lists ◆ ◆ ◆. Meta names Scope3 (now Apostra), which it says acquired Adloox's measurement operations, as a block-list partner; Peer39 has owned Adloox since June 2026 ◆ ◆ ◆. Pixalate documents pre-bid filtering inside Amazon DSP that Amazon configures and runs; Amazon's own pages opened do not name Pixalate ◆.

No screened page documents verifier post-bid blocking on platform-owned placements outside Amazon ◆. Amazon accepts blocking tags without identifying supply, while IAS's letter includes blocked-display metrics for some Amazon sources ◆. Sampling rates are unstated, and only Meta, Snap, IAS-on-Snap and IAS-on-Pinterest document reporting freshness ◆. Billing remains a separate question (Chapter 4).

PLATFORM	THIRD-PARTY PRE-BID CONTROL	POST-BID BLOCKING BY A VERIFIER	EVIDENCE
Meta	Pre-screen and block lists from DoubleVerify, IAS, Zefr and Scope3, applied by Meta	Peer39 describes Adloox blocking; no mechanism stated	◆ ◆ ◆ ◆
YouTube	Not stated for YouTube; Channel Factory's lists, built through YouTube's interface (a pre-bid reading)	None documented	◆ ◆
TikTok	Video exclusion lists managed by DoubleVerify, IAS or Zefr and executed by TikTok, for Standard and Limited inventory; IAS pre-bid in unnamed markets (2022)	None documented	◆ ◆
Amazon	Pixalate's pre-bid invalid-traffic filter, configured and run by Amazon, by Pixalate's account	Open: blocking tags accepted on supply not named; blocked-ad metrics in IAS's letter	◆ ◆

PLATFORM	THIRD-PARTY PRE-BID CONTROL	POST-BID BLOCKING BY A VERIFIER	EVIDENCE
X	IAS pre-bid on the vertical video surface only, for US advertisers who opt in; mechanism not stated	None documented	◆ ◆
Snap	IAS floor classifications sent to Snap, which acts on them	None documented	◆
Pinterest	Not stated	None documented	◆ ◆
Reddit	Limited inventory tier powered by IAS, managed accounts only	None documented	◆
LinkedIn	None on the feed; DoubleVerify pre-bid filtering on the Audience Network, which is off-platform supply	None documented	◆

The platforms' own controls and numbers

Each platform also offers its own controls. Among the nine, only Google's YouTube Inventory Modes hold MRC content-level accreditation, for in-stream video and Shorts; Meta's was revoked in September 2025 on its withdrawal, and the other platforms' tiers, such as TikTok's Inventory Filter, Snap's Brand Suitability Suite and X's Sensitivity Settings, are unaccredited ◆. TikTok calls its invalid-traffic handling certified under the Trustworthy Accountability Group's (TAG) Certified Against Fraud programme and powered by Protected by Mediaocean (formerly Protected Media), whose MRC accreditation is the vendor's own and does not extend to TikTok; the page does not say which TikTok traffic the vendor filters ◆ ◆. Which framework edition each tier follows is Chapter 8's question.

The table's pass rates come from inventory sellers and apply to the content and impressions made measurable ◆ ◆. IAS defines its denominator in two ways, all served ads and measured impressions ◆. Figures invoking GARM use a framework whose owner announced discontinuation on 9 August 2024 ◆.

PLATFORM	WHAT THE PLATFORM STATES	BASE AND PERIOD, AS STATED	MEASURED BY, PER THE PLATFORM	EVIDENCE
X	99.99% brand safety and 97% brand suitability, as averages	2025; base not stated	IAS and DoubleVerify	◆
X	More than 99% compliance with industry-defined brand-safety floors	Period and base not stated	IAS and DoubleVerify	[97]
X	More than 99% efficacy of its keyword and account adjacency controls	Period and base not stated	X itself	[98]
Snap	99% of Spotlight and creator content brand safe against the GARM floor	Four beta advertisers, two weeks of live campaigns, US and EU, May 2024	IAS	[99]
Snap	99% of Spotlight and Creator Stories content brand safe on average	30 days before 24 February 2025; base not stated	IAS	[100]
Meta	Under 1% of Facebook Feed content in GARM's high-risk category	Early testing, before 30 March 2023; base not stated	Zefr	◆
Meta	More than 99% of content next to ads in Feed and Reels brand safe	30 days before 13 February 2025; brand safety not defined	Meta's brand safety and suitability partners	◆
TikTok	More than 26 billion impressions measured, in over 20 languages	The year before 19 May 2023	DoubleVerify, IAS and Zefr	◆
TikTok	Misinformation adjacency below 0.1%	May to August 2024	Zefr	◆

H4 under its rule

H4's contradiction rule requires one documented verifier-executed route on a platform's own placements ◆. TikTok meets the literal test on DoubleVerify's 2021 and IAS's 2022 OM SDK releases for viewability and IVT ◆. Confidence is medium because the current route is not restated and TikTok controls the integration ◆. A platform-control interpretation would change that coding, so it remains an explicit alternative for the author ◆.

On the other eight platforms, the pages screened document no verifier code running in the platform's own placements, Amazon's tags on unnamed supply aside, and no verifier crawl of the content next to ads ◆. Platform-supplied routes are documented for Meta's feeds and in-stream video, for YouTube and for Amazon; Snap and LinkedIn document trackers without saying where they fire; X, Pinterest and Reddit name partners without a route, which the rule codes unresolved ◆. Outside TikTok, the brand-safety and suitability route is platform-supplied where documented (Meta, YouTube's exposure data and Amazon) and unresolved on X, Snap, Pinterest, Reddit and LinkedIn's own feed ◆. One platform of the nine has a documented verifier-executed route, on documents from 2021 and 2022, and every route here is due to be rechecked at the freeze on 11 December 2026 ◆ ◆.

CHAPTER 6

CTV and in-app

On MRC's list read on 10 November 2026, five scored products hold CTV SIVT accreditation, two hold certified-app video viewability, and Zefr holds English-language YouTube in-stream content-level accreditation ◆ ◆ ◆. Streaming programme signals outside YouTube lack content-level accreditation in that screen ◆. TV-off detection is not mandatory, and OM SDK adoption is not measured by its certified-integration count ◆ ◆. The fraud-scheme disclosures below come from detection sellers and publish no error rates ◆.

IN SHORT

Of 41 MRC services in the 10 November 2026 screen,18 carry a CTV metric and eight hold CTV SIVT filtration ◆. Five SIVT holders are scored products, while only DoubleVerify and IAS hold certified-app CTV video viewability ◆ ◆ ◆. Zefr's English-language YouTube in-stream reporting is the sole CTV content-level grant ◆. The preliminary six-product tally of 14 accredited units among 17 requires reconciliation with the full sold-unit grid ◆. MRC's 2021 guidance makes TV-off detection optional and requires unknown treatment for SSAI traffic that cannot be discretely evaluated for IVT ◆ ◆. The OM SDK register held 26 active CTV integrations from 12 companies, giving no impression-level adoption share ◆. All 15 schemes below were disclosed by detection sellers, with no error rate in the 14 disclosures opened ◆. The BADBOX family also has government and court records, whose allegations and procedural stages remain separate ◆.

NOTE ON A DECLARED INTEREST

The author works at Samba TV, which competes with Nielsen in TV measurement. DoubleVerify agreed on 6 August 2026 to be acquired by Neptune BidCo US Inc., Nielsen's parent, and no closing was on record on 10 November 2026 (Chapter 7) ◆ ◆. DoubleVerify is treated under the same published rules as every other vendor in this paper.

What MRC accredits on CTV

CTV SIVT is the widest-accredited family in this screen, with eight holders, five scored ◆. MRC treats CTV as a separate environment and does not accredit OTT discretely ◆. Google Ads and DV360 are first-party validation, while other holders include ad servers, publishers and measurement services ◆. DoubleVerify estimates CTV at 12% of its 2025 post-campaign measurement ◆. Chapter 2 explains what the letters certify.

DoubleVerify's 13 April 2026 letter covers CTV impressions with SIVT, video viewability, property-level verification in ten languages and specified pre-bid segments ◆. Certified apps limit viewable impressions and Fully On-Screen, while other tested apps may enter other viewability metrics ◆. Pre-bid suitability scope remains narrower than its 2024 release might imply (Chapter 2) ◆ ◆.

IAS's letter of 9 July 2026 accredits CTV video tracked ads, impressions, viewability and SIVT filtration for certified applications only, and lists its counts and rates of general invalid traffic (GIVT) and SIVT for desktop, mobile web and in-app only ◆. Its CTV invalid-traffic accreditation is therefore filtration applied to certified-application metrics, first granted in April 2024 ◆ ◆. IAS's ad-fraud page drops the certified-applications limit and claims a priority on CTV SIVT that MRC's archived lists contradict (Chapter 2) ◆ ◆.

Protected Media's and Pixalate's CTV accreditations cover impressions and invalid traffic, not viewability, and so does HUMAN's ◆ ◆. HUMAN holds CTV detection before and after the bid, and no HUMAN CTV viewability product was found in the pages read ◆ ◆. Pixalate's letter of 8 April 2025 is still the one MRC links, and Pixalate's documentation says it reports no CTV viewability ◆ ◆. Protected by Mediaocean, which MRC still lists as Protected Media, holds CTV accreditation for video impressions only, with GIVT and SIVT metrics, by what its letter calls non-SDK measurement ◆. Pre-bid detection of invalid traffic is accredited on CTV for DoubleVerify and HUMAN only, and MRC does not require pre-bid filtration ◆ ◆. Adloox, owned by Peer39 since June 2026, has no CTV row ◆.

In-app accreditation is broader. DoubleVerify, IAS and HUMAN hold impressions with SIVT plus display and video viewability ◆ ◆ ◆. Pixalate holds impressions and display viewability, and Adloox's in-app metrics are net of GIVT only ◆ ◆.

On the large platforms' own CTV inventory, third-party accreditation is thinner (Chapter 5). The verifiers' accredited YouTube viewability covers desktop, mobile web and in-app only, and IAS's accredited reporting for Amazon's demand-side platform excludes Fire TV and Twitch ◆ ◆. MRC's platform chart of 22 October 2026 shows third-party viewability neither accredited nor in process on Fire TV, Twitch and Prime Video ◆.

SERVICE ON MRC'S LIST	CTV METRIC FAMILY	STATUS ON 10 NOVEMBER 2026	SCOPE LIMIT OR NOTE	LETTER	EVIDENCE
DoubleVerify Quality Analytics Platform	Display and video impressions	Accredited	Inclusive of SIVT filtration	13 Apr 2026	◆
DoubleVerify	GIVT and SIVT filtration	Accredited	Trustworthy Accountability Group (TAG) addendum: IVT 2.0 GIVT and SIVT provisions for CTV metrics	13 Apr 2026	◆ [101]
DoubleVerify	Video viewability, Fully On-Screen, completion and quartiles	Accredited	Viewable impressions and Fully On-Screen from certified apps only; other viewability metrics may include tested, non-certified apps; completion and quartiles need no app certification	13 Apr 2026	◆ [101]
DoubleVerify	Property-level ad verification	Accredited	Ten CTV languages (Chinese, English, French, German, Hindi, Japanese, Korean, Malay, Spanish, Vietnamese; French added April 2026); DoubleVerify's own page says nine	13 Apr 2026	◆

SERVICE ON MRC'S LIST	CTV METRIC FAMILY	STATUS ON 10 NOVEMBER 2026	SCOPE LIMIT OR NOTE	LETTER	EVIDENCE
DoubleVerify	Pre-bid segments	Accredited for fraud and IVT, Fully On-Screen, app-store categories and site and app lists	Pre-bid brand safety, suitability and contextual segments on desktop, mobile web and in-app only	13 Apr 2026	◆
DoubleVerify	Content-level brand safety and suitability	Not accredited, not under review	None	None	◆ ◆
Integral Ad Science	Video tracked ads, impressions and viewability	Accredited	Certified applications only; video only; the enclosure has no CTV tracked-ads or Total Net Impressions row; the 2024 letter named five operating systems or devices	9 Jul 2026	◆ ◆
Integral Ad Science	GIVT and SIVT filtration	Accredited; CTV SIVT first granted April 2024	Certified applications only; counts and rates listed for desktop, mobile web and in-app only; standards paragraph names only IVT 2.0's general provisions, the TAG addendum general and sophisticated	9 Jul 2026	◆ [67]
Integral Ad Science	SSAI Noncompliant Ads	Accredited	Certified traffic only (list note)	9 Jul 2026	◆ [47]
Integral Ad Science	Property-level ad verification	No CTV row	Desktop and mobile web on the list; letter basis unresolved	None	◆
IAS YouTube Brand Safety	Content-level brand safety and suitability, English and Spanish YouTube user-generated video	Audit listed as in process as of April 2023	Desktop, mobile web, in-app and CTV	None	◆
HUMAN	Tracked ads and impressions (Ad Fraud Sensor)	Accredited	None stated	14 Apr 2026	◆
HUMAN	GIVT and SIVT detection after the bid; GIVT and SIVT requests before it (Ad Fraud Defense)	Accredited	List: SIVT applies to backend detection and IVT predictions; MRC (April 2025): pre-bid suggestions go to platforms, which make the bid decisions	14 Apr 2026	◆ ◆
HUMAN	Viewability	Not accredited on CTV	Granted on desktop, mobile web and in-app	14 Apr 2026	◆
Pixalate	Display and video ad counts and impressions	Accredited	The letter includes CTV display impressions; the list omits them	8 Apr 2025	◆
Pixalate	GIVT and SIVT filtration	Accredited	Standards: MRC's Minimum Standards and IVT 2.0's GIVT and SIVT provisions; no later Pixalate decision in MRC's quarterly reports, third quarter of 2025 to third quarter of 2026	8 Apr 2025	◆ ◆
Pixalate	SSAI video metrics	Accredited	None stated	8 Apr 2025	◆
Pixalate	Viewability	Neither accredited nor under review on CTV	Video viewability audit listed as in process as of September 2025, for desktop, mobile web and in-app only	None	◆
Protected Media (Protected by Mediaocean)	Video impressions	Accredited	Non-SDK measurement; inclusive of SIVT filtration	22 Sep 2026	◆
Protected Media	Video GIVT and SIVT metrics	Accredited	None stated	22 Sep 2026	◆
Protected Media	Viewability	Not accredited on CTV	None	None	◆

SERVICE ON MRC'S LIST	CTV METRIC FAMILY	STATUS ON 10 NOVEMBER 2026	SCOPE LIMIT OR NOTE	LETTER	EVIDENCE
Zefr Third-Party Reporting of Google YouTube Content-Level Brand Safety	Content-level brand safety and suitability	Accredited, initial grant	English-language in-stream video and music video; Shorts, YouTube TV, Google TV, Masthead, Google Video Partner videos and volumetric data excluded	20 Jan 2026	◆
Adloox (Peer39)	Any CTV metric	No CTV row	None	None	◆
Google YouTube Brand Safety (first-party, not scored)	Content-level brand safety and suitability	No CTV on the list	Desktop, mobile web and in-app	28 May 2026	◆ [68]
GumGum Contextual Platform API (listed, not scored)	Content-level brand safety, open web	Revoked September 2025 on withdrawal	English; had covered desktop, mobile web and CTV	Continued May 2025	◆

Content-level and program-level suitability

Zefr's 20 January 2026 grant covers English-language YouTube in-stream video and music video on CTV ◆. It excludes Shorts, YouTube TV, Google TV, Masthead, Google Video Partners and volumetric data ◆. Google's first-party row lists no CTV, IAS's YouTube audit remains in process, and GumGum's former open-web grant was revoked on withdrawal in September 2025 ◆.

Vendors sell or describe program-level signals for streaming without that accreditation ◆. On 6 November 2025 DoubleVerify launched Verified Streaming TV, a check of premium streaming-player environments, and Do Not Air automation, which curates streaming inventory by genre and maturity rating and was offered in early access through The Trade Desk ◆. Pixalate's OpenEPG resolves the show in a CTV bid request, and IAS sells CTV Brand Safety for YouTube while its YouTube audit is listed as in process ◆. Channel Factory added genre and mood classification for premium CTV on 28 September 2026 ◆. Peer39 says roughly 60% of CTV bid requests lack usable program-level signal, with no sample, period or method ◆.

Program metadata often comes from the seller: IAB Tech Lab notes that genre travels as free text in OpenRTB bid requests, and DoubleVerify's Certified Transparent Streaming programme takes content data that publishers contribute through clean rooms [102, 103]. Classifying a program from its genre, rating or title differs from the content-level measurement MRC accredits, which no CTV service outside YouTube held on 10 November 2026 ◆ ◆.

CTV units in the final grid

The final H1a grid (Chapter 2) records 34 CTV units sold by 14 of the 17 scored products, of which 13 are MRC-accredited and none is in audit ◆. The grid counts one unit per product, metric family and environment, so pre-bid and post-bid delivery of one family on CTV is one unit. YouTube is one of the nine large platforms and has its own units, so YouTube served on a TV screen, such as Zefr's content-level accreditation, is counted under YouTube, not CTV. The stream's earlier six-product working tally (14 of 17) is superseded by this grid ◆.

Server-side insertion and the switched-off screen

SSAI ad-stitching servers send tracking requests on devices' behalf ◆. VAST4.3 requires device IP and user-agent forwarding to distinguish them from fraud, but White Ops says these fields are easy to falsify ◆. MRC's guidance governs their validity treatment below ◆.

MRC's SSAI and OTT guidance of August 2021 makes vendors verify that providers declaring SSAI are SSAI providers, and treat their traffic as unknown rather than valid unless it is discretely measurable for IVT; proxy and data-centre traffic may be filtered only when direct signals show it is invalid ◆. The guidance says CTV may be more vulnerable to ad fraud, partly because measurers' telemetry there may be more limited ◆. It encourages ads.cert, IAB Tech Lab's protocol for signing requests between servers, and proposes an SSAI-provider certification that no MRC or Tech Lab page searched on 10 November 2026 shows established ◆ ◆. MRC's IVT 2.0 addendum (June 2020) has SIVT vendors consider in-app and OTT traffic discretely where it is a material share of what they measure, and its April 2024 interim update says non-CTV apps misrepresented as CTV should be considered SIVT ◆.

Pixalate's SSAI video metrics and IAS's count of SSAI Noncompliant Ads on CTV are accredited ◆ ◆. Pixalate says it judges SSAI traffic per transaction from more than 30 signals about the SSAI server, and publishes no signal list, threshold or validation result ◆. TAG's Certified Against Fraud 11.0 (July 2026) dropped its CTV SSAI billing-notification requirement ◆.

MRC's 2021 guidance does not require TV-off detection, but requires removal of confirmed TV-off events and disclosure of the limitation ◆. Unflagged plays on a switched-off screen can remain in an accredited count ◆. DoubleVerify's certified-app Fully On-Screen metric adds its TV-on condition ◆ [103]. Its 2024 estimate of more than one in three impressions in TV-off-capable environments had no published method or base ◆.

Open Measurement on CTV and in-app

On CTV, what is public about the OM SDK is its releases and its certified integrations ◆ ◆. The kit lets a measurement provider's script collect signals inside an app or player that integrates it ◆ ◆. Version 1.4 (2022) first supported CTV, on Android TV and Apple's tvOS, with a signal for a display that is off ◆. Version 1.5 (announced 30 May 2024) added Samsung and LG smart-TV web apps, and device attestation for Fire TV and Apple devices, announced on 23 October 2025, reached the JavaScript client as 1.6.0 on 29 October ◆ [104]. On 10 November 2026 IAB Tech Lab announced impression measurement for pause, menu, screensaver and overlay ads on CTV, naming no version [105]. No Tech Lab page opened names Roku or Vizio support ◆.

On 10 November 2026 IAB Tech Lab's compliance register listed 341 active certified integrations from 134 companies, plus 125 expired records; most active records are for Android and iOS apps, and 26 active records from 12 companies carried a CTV label ◆. The records are mostly SDKs, players and ad-insertion services, with some individual apps ◆. Thirty-two active records were last certified before 2023, and DIRECTV's CTV-Android record, which a May 2024 Tech Lab release named, shows as expired ◆. That release's unsourced claim that the OM SDK covered 40% of the CTV market measures platform

support, not adoption ◆. Companies can hold records on several platforms, so the rows below do not sum to 134 [106].

PLATFORM LABEL	ACTIVE RECORDS	COMPANIES	V1.2	V1.3	V1.4	V1.5	V1.6	EVIDENCE
Android	145	127	4	17	33	27	64	[106]
iOS	136	120	3	19	18	30	66	[106]
Web	34	30	0	3	4	17	10	[106]
CTV-Android	11	10	0	0	3	6	2	[106]
CTV-iOS (tvOS)	10	9	0	0	2	3	5	[106]
CTV-Web (Samsung, LG)	4	4	0	0	0	2	2	[106]
CTV, platform not specified	1	1	0	0	0	0	1	[106]
All active	341	134	7	39	60	85	150	◆
All active CTV	26	12	0	0	5	11	10	◆

MRC's letters tie OM SDK measurement to DoubleVerify's non-CTV viewability and IAS's in-app measurement, with MRAID metrics excluded ◆. DoubleVerify separately claimed CTV OM SDK support in December 2025 ◆. TAG's 2026 rules require verifiers under the device-attestation guidance to use 1.6 or higher, while impression-level adoption remains unreported in the screened sources ◆ ◆.

The answers to device spoofing are proprietary or unmeasured ◆. App-ads.txt extends authorised-seller files to CTV app stores but authenticates no device, and IAB Tech Lab's CTV Programmatic Guide, released for public comment in December 2025, calls device information in bid requests a text string that is easily manipulated ◆ [107]. Roku's Advertising Watermark is designed to let only genuine Roku devices display authenticated ads, and OM SDK attestation covers Apple devices and Fire TV ◆ [104]. The ads.cert protocol for authenticating devices is still marked future, its repository's last commit is dated 8 March 2023, and no adoption figure was found in the sources opened ◆. IAB Tech Lab's release names DoubleVerify, HUMAN and IAS, which sell detection, among the companies in the group behind attestation [104]. No source opened reports how much CTV inventory carries any of these signals ◆.

Schemes and who found them

Each scheme in the register is its finder's finding, and every finder sold the detection it publicised at the time ◆. Of the 15 schemes, disclosed from 2020 to 2026, White Ops and HUMAN named nine and DoubleVerify four; Oracle, then owner of the Moat verification service, and Pixalate named one each ◆. Every figure in the table is the finder's, in its own unit, and the units cannot be added ◆.

SCHEME	FINDER, AND WHAT IT SOLD	DISCLOSED	ENVIRONMENT AND TECHNIQUE	SCALE AS THE FINDER STATED IT	EVIDENCE
Monarch	Pixalate (IVT analytics and blocking)	25 Mar 2020	CTV (Roku apps): app spoofing and falsified measurement events	Hundreds of thousands of IVT impressions a day for over six months across over a dozen apps, on "conservative estimates"; asserts no culpability	◆

SCHEME	FINDER, AND WHAT IT SOLD	DISCLOSED	ENVIRONMENT AND TECHNIQUE	SCALE AS THE FINDER STATED IT	EVIDENCE
ICEBUCKET	White Ops, now HUMAN (bot mitigation and pre-bid blocking)	16 Apr 2020	CTV and mobile SSAI: data-centre traffic with forged device details	About 1.9 billion ad requests a day at the January 2020 peak; nearly 28% of the programmatic CTV traffic White Ops could see and 66% of the programmatic CTV-related SSAI traffic it protected that month; more than 2 million people impersonated	◆
StreamScam	Oracle (Moat verification, retired 30 September 2024)	17 Dec 2020	CTV SSAI: servers sent forged impression events	More than 28.8 million US household IP addresses; about 3,600 apps; 3,400 CTV device models	◆ ◆
CopyCatz	White Ops (bot and ad-fraud detection)	Jan 2021	Mobile in-app (Android): copycat apps	164 apps; more than 10 million downloads	◆
OctoBot, with MultiTerra and SneakyTerra	DoubleVerify (fraud protection)	20 Apr 2021	CTV: seven related spoofing variants since November 2019	MultiTerra about $1 million and SneakyTerra about $5 million a month, method not stated; not re-opened in the verifier pass	◆
PARETO	HUMAN (ad-fraud detection)	21 Apr 2021	CTV, spoofed from infected Android devices	Nearly 1 million infected mobile Android devices; more than 6,000 CTV apps spoofed; an average 650 million ad requests a day	◆
SmokeScreen	DoubleVerify (fraud protection)	12 Aug 2021	CTV screensaver apps, including with the TV off	About 10,000 devices and up to 10 million requests a day; over $6 million a month by pricing ad requests at $20 a thousand; not re-opened in the verifier pass	◆
LeoTerra variant	DoubleVerify (fraud protection)	24 Aug 2022	CTV SSAI: spoofed device lists, including connected appliances	Up to $10 million in 2022; more than 92 million devices spoofed in the first half; not re-opened in the verifier pass	◆
Scylla	HUMAN (ad-fraud detection)	2022, restated Jan 2023	Mobile in-app (9 iOS and 80 Android apps)	89 apps; more than 13 million downloads; known only from the VASTFLUX release	◆
VASTFLUX	HUMAN (ad-fraud detection)	19 Jan 2023	Mobile in-app, mainly iOS: stacked, invisible video ads	More than 1,700 apps spoofed; nearly 11 million devices; a peak of 12 billion ad requests a day	◆
BADBOX and PEACHPIT	HUMAN (ad-fraud detection)	4 Oct 2023	Backdoored CTV boxes, phones and tablets; hidden ads	More than 74,000 infected devices; PEACHPIT's 39 apps averaging 4 billion ad requests a day	◆ ◆
CycloneBot	DoubleVerify, with Roku (fraud protection)	28 Feb 2024	CTV: bots spoofing sessions and devices	Capable of up to 250 million falsified ad requests and about 1.5 million spoofed devices a day; up to $7.5 million a month on industry price estimates	◆
Konfety	HUMAN (ad-fraud detection)	16 Jul 2024	Mobile in-app (Android): evil-twin apps	More than 250 decoy apps; a peak of 10 billion bid requests a day	◆
BADBOX 2.0	HUMAN, with Google, Trend Micro and Shadowserver (ad-fraud detection)	5 Mar 2025	Backdoored CTV boxes and other Android devices	More than 1 million infected devices; called by HUMAN the largest CTV botnet ever disclosed	◆
NewsJunkie	HUMAN (ad-fraud detection)	7 Jul 2026	CTV: device spoofing	Hundreds of millions to nearly 2 billion invalid CTV bid requests a day per seller at peak; number of sellers and period not given	◆ ◆

None of the 14 disclosures opened reports a miss rate, a false-positive rate or precision, and 12 say the finder's customers were protected or its product detected the scheme ◆. The research was done by AI agents under the author's direction; in the verifier pass a separate

AI agent, told to refute each finding, re-opened ten of the 14 (Chapter 13) ◆. The OctoBot, SmokeScreen and LeoTerra figures rest on the research agents' reads alone ◆.

Several schemes attacked measurement itself ◆. Oracle said StreamScam's servers sent Moat and advertisers impression events with forged IP addresses, app IDs and device IDs, and delivered no ad ◆. ICEBUCKET hid bots in the limited signal of SSAI traffic, and VASTFLUX stacked invisible video ads where iOS apps give verification partners little signal ◆.

Scale figures travel beyond their base. ICEBUCKET's 28% was a share of the programmatic CTV traffic White Ops could see, but the briefing TAG co-hosted on 30 April 2020 captioned it as global CTV traffic affected ◆. Roku said PARETO affected less than 0.1% of its devices and that CycloneBot did not affect its platform [108, 109].

Only the BADBOX family has government and court records, and they draw on the finders ◆. A Federal Bureau of Investigation (FBI) announcement of 5 June 2025 counts millions of infected devices without mentioning advertising ◆. Google's complaint in Google LLC v. Does 1-25 (Southern District of New York, 1:25-cv-04503) alleges more than ten million infected devices as of April 2025, against HUMAN's count of more than one million in March 2025 ◆ ◆. A complaint is an allegation, and the preliminary injunction of 27 June 2025 is not a judgment ◆. Google's statement that it refunds advertisers for fraudulent traffic describes first-party validation (Chapter 4) ◆.

DoubleVerify reports 140% more detected CTV schemes and variants in Q1 2026 than Q1 2025, without raw counts or a detection-validation method ◆. The period is stated, so the figure is a vendor trend within that period ◆. None of this chapter's screened sources independently measures a verifier's CTV error rate ◆. Chapter 3 covers prevalence comparisons.

Standards and specifications, 2019 to 2026

Each row is dated from the issuer's own document as read on 10 November 2026. A guidance date is not an accreditation date, and a release is not deployment ◆.

DATE	ISSUER	DOCUMENT	WHAT IT SAYS FOR CTV AND IN-APP VERIFICATION	EVIDENCE
March 2019	IAB Tech Lab	app-ads.txt (ads.txt 1.0.2)	Authorised-seller files for apps, including connected-TV app stores	◆ [110]
June 2020	MRC	IVT 2.0 addendum, section 4.2.2	In-app and OTT traffic considered discretely where material; OTT proxy and data-centre traffic filtered only when known to be invalid	◆
March 2021	IAB Tech Lab	ads.txt 1.0.3	INVENTORYDOMAIN variable for inventory shared between sellers	[111]
August 2021	MRC	Server-Side Ad Insertion and OTT Guidance	CTV accredited discretely, OTT not; TV-off detection not required; unverifiable SSAI traffic unknown; ads.cert encouraged	◆ ◆
2022	IAB Tech Lab	OM SDK 1.4	First CTV release: Android TV and tvOS; display-connection status	◆

DATE	ISSUER	DOCUMENT	WHAT IT SAYS FOR CTV AND IN-APP VERIFICATION	EVIDENCE
August 2022	IAB Tech Lab	ads.txt 1.1	OWNERDOMAIN and MANAGERDOMAIN variables	[111]
December 2022	IAB Tech Lab	VAST 4.3	Device IP address and user agent sent with ad-stitching tracking requests	◆
14 February 2024	IAB Tech Lab	ads.cert 2.0 page, last updated	Authenticated Connections ready for adoption; Authenticated Devices future	◆
24 April 2024	MRC	2024 IVT Interim Updates	CTV bundle IDs within app misrepresentation; non-CTV apps posing as CTV are SIVT; app-ads.txt telemetry encouraged	◆
30 May 2024	IAB Tech Lab	OM SDK 1.5 announced	Samsung Tizen and LG webOS through the Web Video SDK	◆
2024	IAB Tech Lab	VAST CTV Addendum 2024	AdVerification and UniversalAdId fields extended to VAST 2 and 3	[112]
11 April 2025	MRC	Statement on pre-bid IVT	Post-serve GIVT filtration required on every impression; pre-bid filtration not required	◆
23 and 29 October 2025	IAB Tech Lab	Device attestation announced; JavaScript client 1.6.0	Privacy Pass attestation for Apple devices and Fire TV	◆ [104]
11 December 2025	IAB Tech Lab	CTV Programmatic Guide, for public comment to 31 January 2026; page modified 30 January 2026	Use cases for OM SDK, app-ads.txt, ads.cert, device attestation and content taxonomy	◆ [107]
July 2026	TAG	Certified Against Fraud 11.0	CTV SSAI billing-notification requirement removed; verifiers under its device-attestation guidance must implement OM SDK 1.6 or higher	◆ ◆
2 November 2026	IAB Tech Lab	OM SDK JavaScript client 1.6.12	Version alignment; no change from 1.6.11	◆
10 November 2026	IAB Tech Lab	OM SDK support for emerging CTV ad formats	Impression measurement for pause, menu, screensaver and overlay ads; no version named	[105]

CHAPTER 7

Who pays the verifiers, and who owns them

Advertisers pay most of the verification bill, and sellers pay part of it. In FY2024, the last fiscal year both verifiers filed, DoubleVerify's Supply-side customers brought it 8.6% of its revenue, and the Publisher channel of Integral Ad Science (IAS), which includes Publica's connected-TV (CTV) ad serving, brought IAS 14.4%. Within eight months both companies changed hands or agreed to: Novacap's funds took IAS private on 23 December 2025, and DoubleVerify agreed on 6 August 2026 to be acquired by Neptune BidCo US Inc., Nielsen's parent, a sale still pending on 10 November 2026. H5, the prediction that no verifier with seller revenue or an interested owner publishes a safeguard, is contradicted, but narrowly, by one DoubleVerify page that describes a structural separation for its publisher clients and says nothing of the pending owner. Ownership is recorded as a governance attribute, never as evidence that a measurement is biased.

IN SHORT

In FY2024, DoubleVerify's Supply-side customers paid $56.8 million of $656.8 million revenue (8.6%), while IAS's Publisher channel, including Publica's CTV ad serving, brought $76.5 million of $530.1 million (14.4%) ◆ ◆ ◆. IAS's last periodic filing at the 10 November 2026 check was its 10-Q for the quarter to 30 September 2025 ◆. DoubleVerify's measured transactions rose from about 4.5 trillion in 2021 to 9.5 trillion in 2025, while its average fee fell 4% in FY2024 and 3% in FY2025, excluding one introductory deal ◆ ◆. Platforms collect or facilitate substantial revenue for both vendors, on different measures ◆. IAS closed its take-private on 23 December 2025; DoubleVerify's agreed sale to Neptune BidCo US Inc., Nielsen's parent, remained pending on 10 November 2026 ◆ ◆. H5 is contradicted on DoubleVerify's June 2024 separation statement for publisher clients, although its scope, responsible owner and enforcement are unspecified and it does not address the pending owner ◆ ◆. A stricter policy requirement would yield a different reading within the sources screened, as set out below ◆.

Advertisers pay most of the bill

Advertisers and their agencies pay most of what the two verifiers with public filings earn; sellers pay a steady minority share. DoubleVerify files revenue by three customer types: Activation, from advertisers that evaluate inventory before buying; Measurement, from

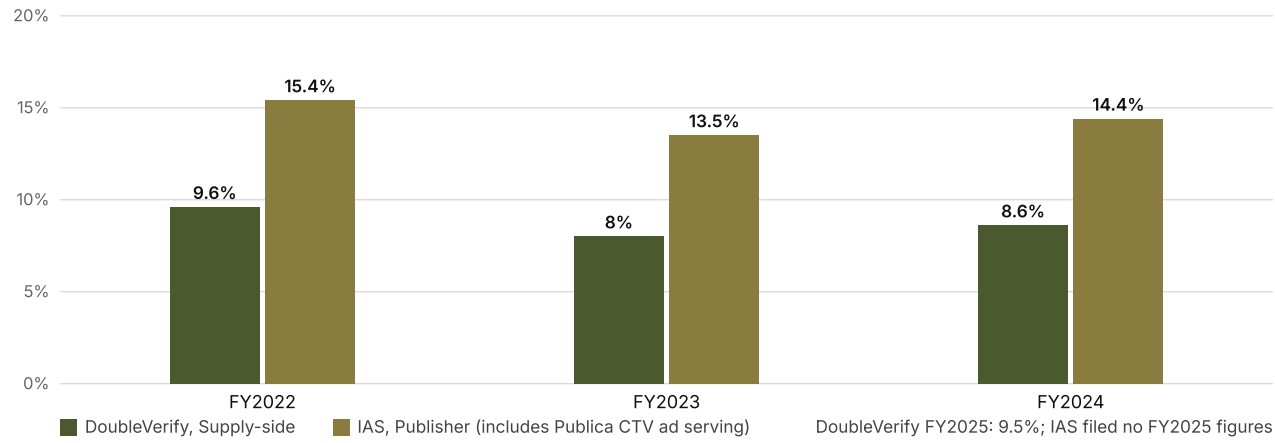
advertisers that measure quality and performance afterwards; and Supply-side, from publishers, retail media networks, ad networks and exchanges [26]. In FY2025 the three brought $427.3 million, $249.7 million and $71.3 million of $748.3 million. Its annual report, the 10-K, rounds the advertiser share to 90%, against 91% in FY2024 [26]. Supply-side customers sign monthly or annual contracts with minimum guarantees and tiered pricing above them [26].

IAS files revenue by channel: Optimization, Measurement and Publisher, which in FY2024 brought $242.6 million, $211.0 million and $76.5 million of $530.1 million; 86% of revenue came from advertiser customers. The Publisher channel mixes tools sold to publishers with the CTV ad serving of Publica, which IAS bought in August 2021, and IAS credited most of the channel's 2024 growth, $12.7 million or 20%, to Publica [27]. Neither company splits its seller-side line into verification and other products, so the ad-serving part of IAS's line is unknown.

SELLER-SIDE LINE, AS FILED	FY2022	FY2023	FY2024	FY2025	EVIDENCE
DoubleVerify, Supply-side (customer type)	$43.3m of $452.4m (9.6%)	$45.6m of $572.5m (8.0%)	$56.8m of $656.8m (8.6%)	$71.3m of $748.3m (9.5%)	
IAS, Publisher (channel; includes Publica's CTV ad serving)	$62.8m of $408.3m (15.4%)	$63.8m of $474.4m (13.5%)	$76.5m of $530.1m (14.4%)	Not filed	

The comparison uses FY2024, the last year both companies filed, while keeping the customer-type and channel definitions apart. DoubleVerify's Supply-side share was 9.8% in the nine months to 30 September 2026, $56.8 million of $580.3 million. IAS's last filed period was the nine months to 30 September 2025, with Publisher revenue of $68.7 million of $437.6 million. Its FY2025 amount remained unfiled at the 10 November check and is unresolved under H5. The figure plots each line against its own company's revenue.

SELLER-SIDE REVENUE AS A SHARE OF TOTAL REVENUE, AS FILED



Both verifiers earn from sellers: about a tenth of DoubleVerify's revenue, about a seventh of IAS's. DoubleVerify's Supply-side customer type and IAS's Publisher channel as a share of total revenue in the 10-K filings, FY2022 to FY2024 (DoubleVerify FY2025: 9.5%). The two categories are defined differently: IAS's Publisher line includes Publica's CTV ad serving, and neither company says how much of the line is paid by platforms rather than publishers.

Units: percent of total revenue · Scope: DoubleVerify and IAS 10-K filings, fiscal years as filed · Sources [26] [271] [27]

A revenue category is not a map of who pays for each verdict, and existence of seller revenue is not evidence that a verdict is biased (hypothesis H5).

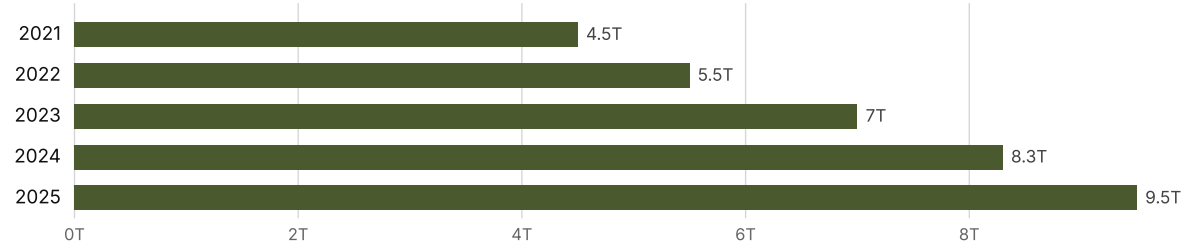
The filed lines do not identify the payer for each verdict. Agencies pay DoubleVerify for multiple clients and platforms collect advertiser fees ◆. DoubleVerify also sells platform-wide measurement subscriptions to supply-side partners [26]. IAS's terms allow a media partner to pay an advertiser's verification fees, without disclosing the amount billed that way ◆. Seller revenue alone establishes no measurement bias.

Volume doubled while the fee per thousand fell

DoubleVerify bills advertisers per thousand Media Transactions Measured, mostly at fixed prices, with caps where platforms charge a share of spend ◆. The approximate volume rose from 4.5 trillion in 2021 to 5.5 trillion in 2022, 7.0 trillion in 2023, 8.3 trillion in 2024 and 9.5 trillion in 2025 ◆. This company-defined operating metric counts measurements before and after purchase, rather than unique impressions or impressions classified as valid; the 2025 mix was an estimated 44% display and 56% video ◆.

Average fees fell 4% in FY2024 and 3% in FY2025 while measured volume rose 19% and 15%, excluding an introductory fixed-fee deal with one unnamed customer ◆. In the nine months to September 2026, fees fell 4% and volume rose 11%; the filings report changes without giving an average dollar fee ◆.

DOUBLEVERIFY MEDIA TRANSACTIONS MEASURED, TRILLIONS A YEAR



The measured-transaction fee per thousand fell 4% in FY2024 (volume +19%) and 3% in FY2025 (volume +15%), excluding an introductory-pricing effect the filings name.

> **DoubleVerify measures twice the volume it measured in 2021 and charges less for each thousand.** Media transactions measured by DoubleVerify, as reported in its 10-K filings, 2021 to 2025, with the change in its measured-transaction fee per thousand. The price of a verdict falls as its volume rises: the fee is set per thousand impressions measured, not by what the verdict changes.
>
> Units: trillions of media transactions measured per year · Scope: DoubleVerify, worldwide, as filed · Sources [26] [271] [272]
>
> A media transaction is DoubleVerify's unit (an impression measured, pre-bid or post-bid); IAS publishes no comparable annual count.

DoubleVerify's partner-pay rate card, effective 21 May 2026 and applying only where a partner remits its fees, lists $0.135 per thousand impressions for DV Authentic Ad and $0.400 for DV Social Pre-Bid Controls; no direct-contract price list was found ◆.

IAS's FY2024 10-K gives no comparable annual measurement count. Its more than 280 billion daily interactions and Publica's more than 10 billion monthly CTV ads use different units and activities ◆. IAS's measurement impressions rose 23% with average costs per thousand (CPMs) down 7%; optimization impressions rose 4% with CPMs up 2% ◆. Revenue follows measured volume and CPMs, with minimums, tiered overages or flat fees in some contracts [27]. The two vendors' volume series cannot be added.

A few platforms collect much of the money

Customer concentration and platform dependence use different measures. No DoubleVerify customer exceeded 10% of revenue in 2023–2025, while IAS's 10-K gives only an accounts-receivable concentration test ◆. Two unnamed platforms collected 24% and 12% of DoubleVerify's FY2025 revenue; DSP partnerships facilitated 46% of IAS's FY2024 revenue, the largest one 18% ◆. These are collected and facilitated revenue, respectively.

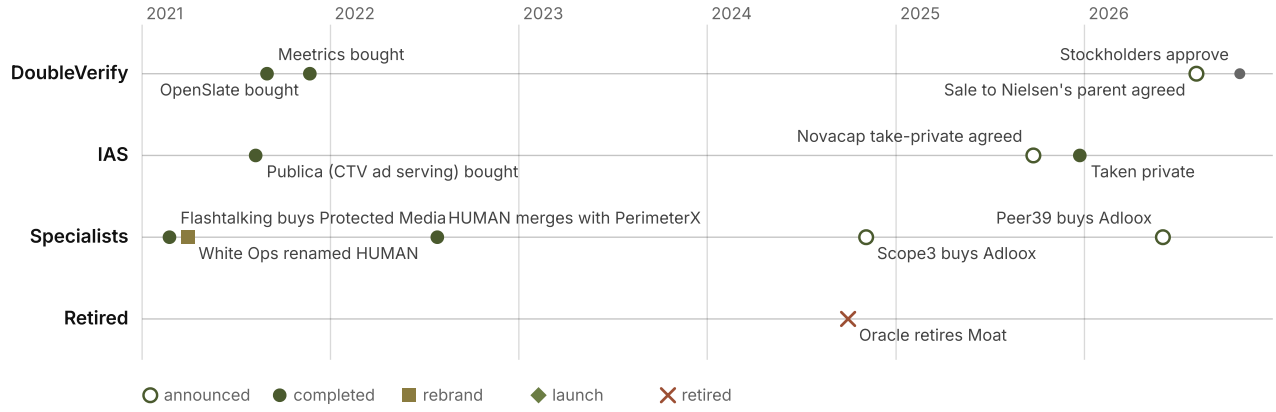
Money also flows back. IAS pays DSPs, including Google and The Trade Desk, a share of the revenue it receives through its products, and says the platforms it integrates with produce no material revenue for it directly ◆ ◆. The platforms that collect the fees are also where pre-bid segments are sold: a dependency, not a finding of conflict ◆.

Large customers are counted in different units. DoubleVerify had 131 customers worth at least $1 million each in 2025, against 93 in 2023 [26]. IAS had 237 advertisers spending at least $200,000 a year in 2024, and they produced 85% of its advertising revenue [27].

Who owns the verifiers now

The dated record includes acquisitions, financings, rebrands and retirements [launch] [launch] [launch] [launch]. DoubleVerify completed its initial public offering on 23 April 2021 at $27 a share; IAS closed its offering on 2 July 2021 at $18 [launch] [launch]. Within eight months, IAS closed a take-private and DoubleVerify agreed to one [launch] [launch].

OWNERSHIP EVENTS IN AD VERIFICATION, 2021 TO 2026



> **Both independent verifiers changed hands or agreed to within eight months.** Dated ownership events from the companies' filings and releases. IAS was taken private on 23 December 2025; DoubleVerify's sale to Neptune BidCo US Inc., Nielsen's parent, was agreed on 6 August 2026 and approved by stockholders on 29 October 2026, and had not closed when last checked (10 November 2026). The full dated table, with value bases, is in the data downloads.
>
> Units: dated events · Scope: verification companies and products, 2021 to 10 November 2026 · Sources [273] [115] [274] [133] [275] [276] [277] [278] [114] [28] [116] [279] [280] [281] [282] [25] [45]
>
> Pending: DoubleVerify's closing is rechecked at the freeze. A headline deal value is shown only with the basis its source states.

Q3, the paper's question on consolidation, is answered under a coding rule: an event is completed only when a filing or the parties' release records completion, and a pending transaction is dated to the latest primary record read. Acquisitions, mergers, financings, rebrands and retirements are different events, and none shows that a product or metric failed [launch] [launch]. The table lists the events other than IAS's take-private and DoubleVerify's sale, which follow.

DATE	COMPANY	EVENT	CLASS	VALUE AS STATED	EVIDENCE
23 Feb 2021	Protected Media	Flashtalking announces it has acquired the company	Acquisition (completed per release)	Not stated	[launch]
30 Mar 2021	White Ops	Renamed HUMAN, three months after Goldman Sachs Merchant Banking Division, ClearSky Security and NightDragon bought it	Rebrand	n/a	[launch]
13 Jul 2021	Flashtalking	Mediaocean agrees to acquire it; no completion date found; part of Mediaocean by 12 May 2022	Acquisition (agreement)	Not disclosed	[launch]
9 Aug 2021	Publica	Acquired by IAS (CTV ad serving)	Acquisition (completed)	$171.4m cash plus shares valued at $49.6m	[launch]

DATE	COMPANY	EVENT	CLASS	VALUE AS STATED	EVIDENCE
31 Aug 2021	Meetrics	Acquired by DoubleVerify (ad verification)	Acquisition (completed)	$24.3m net cash	◆
2 Sep 2021	Pixalate	Growth capital from Western Technology Investment and Javelin Venture Partners	Financing	$18.1m ($22.7m in total)	◆
22 Nov 2021	OpenSlate	Acquired by DoubleVerify (contextual targeting); agreement announced 9 Nov 2021	Acquisition (completed)	$148.2m net of cash, final	◆ [113]
31 Dec 2021	Context (Nobora SAS)	Acquired by IAS (contextual classification)	Acquisition (completed)	$22.6m cash plus shares valued at $10.4m	◆
2021 (day not stated)	Amino Payments	Acquired by IAS	Acquisition (completed)	Not stated	◆
4 Jan 2022	HUMAN	Growth round led by WestCap	Financing	$100m raised	◆
20 Jul 2022	Adverif.ai	Acquired by Zefr (misinformation detection)	Acquisition	Not stated	◆
27 Jul 2022	HUMAN and PerimeterX	Merger under the HUMAN name announced; approvals received; completion not stated	Merger	Not disclosed	◆
14 Aug 2023	Scibids	Acquired by DoubleVerify (campaign optimization)	Acquisition (completed)	$121.4m initially reported; $121.1m net of cash as adjusted in the FY2025 10-K	◆
21 Nov 2023	Kubient (KAI)	Form 25 removes its securities from listing; a current report of 8 May 2024 announces job cuts tied to a potential wind-down	Delisting	n/a	◆
30 Sep 2024	Oracle Moat	Retired with every Oracle Advertising product, after notices in the week of 17 Jun 2024	Product retirement	n/a	◆
9 Oct 2024	HUMAN	Growth capital led by WestCap	Financing	$50m-plus raised	◆
4 Nov 2024	Adloox	Scope3 announces the acquisition	Acquisition (announced as done)	Not disclosed	◆
21 Nov 2024	Innovid	Mediaocean, owner of Protected Media, agrees to acquire it	Acquisition (agreement)	$3.15 a share; no aggregate stated	◆
13 Feb 2025	Innovid	Merger completed; combined with Flashtalking	Acquisition (completed)	$3.15 a share	◆
13 Mar 2025	Rockerbox	Acquired by DoubleVerify (marketing attribution)	Acquisition (completed)	$82.3m net of cash	◆
24 Apr 2025	Channel Factory	Truelink Capital completes a significant investment; Star Mountain exits through it (announced 5 Jun 2025) and stays a minority owner	Investment	Not disclosed	◆
2 Jun 2026	Adloox	Peer39 announces it has bought Adloox from Scope3; completion recorded in a UK filing	Acquisition (completed)	Not disclosed	◆
23 Sep 2026	Scope3	Announces it is becoming Apostra; Adloox's owner unchanged	Rebrand	n/a	◆

DoubleVerify's filings call only one of its four purchases since 2021, Meetrics, an ad verification provider ◆. OpenSlate was a pre-campaign contextual targeting platform which, by DoubleVerify's own announcement, helped brands match ads to suitable or relevant content in social video and CTV; that announcement valued the deal at $150 million, and the final price was $148.2 million net of cash ◆ [113]. Neither brand survives: on 10 November 2026 openslate.com redirected to doubleverify.com and meetrics.com did not resolve ◆.

IAS: agreed, consented, closed

IAS agreed on 24 September 2025 to merge with a subsidiary of Igloo Group Parent, Inc., an affiliate of funds managed by Novacap, for $10.30 a share ◆. Vista and Atlas funds holding about 52.6% consented that day, so the deal used a Schedule 14C information statement [114]. The funding commitments were up to $857.8 million in equity, a $1.0 billion term loan and a $150 million revolver, with termination fees of $52.5 million for IAS and $100 million for the buyer ◆.

IAS became wholly owned by Igloo Group Parent on 23 December 2025 ◆. The closing 8-K reports about $1.6 billion paid for outstanding shares, excluding its defined Owned Company Shares and Appraisal Shares ◆. The parties' earlier $1.9 billion valuation states no basis, and the 19 December event on the 8-K cover concerns retention agreements ◆.

DATE	STEP	RECORD	EVIDENCE
24 Sep 2025	Merger agreement at $10.30 a share; written consent of holders of about 52.6% of the voting power	8-K; joint release	◆
24 Oct 2025	Hart-Scott-Rodino (HSR) antitrust notifications filed; waiting period scheduled to end on 24 Nov 2025 (actual expiry not documented in the sources opened)	Information statement	[114]
7 Nov 2025	Schedule 14C information statement dated and first mailed	Schedule 14C	◆
23 Dec 2025	Merger completed, about $1.6 billion paid for the outstanding shares; Nasdaq files Form 25	Closing 8-K; Form 25	◆ ◆
2 Jan 2026	Form 15 filed, certifying one holder of record	Form 15	◆

DoubleVerify: agreed, cleared, approved, not closed

NOTE ON A DECLARED INTEREST

The author works at Samba TV, which competes with Nielsen in TV measurement. Neptune BidCo US Inc., the parent of Nielsen, is DoubleVerify's pending buyer. DoubleVerify is treated under the same published rules as every other vendor in this paper.

On 6 August 2026 DoubleVerify agreed to become a subsidiary of Neptune BidCo US Inc. at $13.60 a share in cash ◆. The definitive proxy identifies Neptune as Nielsen Company (US), LLC's parent and an affiliate of Elliott- and Brookfield-managed funds ◆. The table separates the completed approvals from the uncompleted merger at the 10 November check ◆ ◆.

The two value figures rest on different bases: an enterprise value of about $2.15 billion in the joint release, an implied equity value of about $2.332 billion in the proxy ◆. Elliott's affiliates committed $200 million of equity and lenders about $1.8 billion of debt; the termination fees are $60 million for DoubleVerify and $144 million for the parent; and Providence's funds, holding about 11.8%, agreed to vote in favour ◆.

DATE	STEP	RECORD	EVIDENCE
6 Aug 2026	Merger agreement with Neptune BidCo US Inc. and Wallace Merger Sub Inc.; the joint release forecasts closing by the end of the fourth quarter of 2026	8-K, Item 1.01 (filed 10 Aug); joint release	◆ ◆
20 Aug 2026	HSR notification filed	Definitive proxy	◆
21 Sep 2026	HSR waiting period expired at 11:59 p.m. ET; Australian notification lodged	Definitive proxy; 8-K of 20 Oct	◆
29 Sep 2026	Definitive proxy; targets closing by the first quarter of 2027 (forecast)	Schedule 14A	◆
5 to 13 Oct 2026	Clearances in Germany (5 Oct), Cyprus (7 Oct), Australia, by the Australian Competition and Consumer Commission (ACCC), and Morocco (13 Oct)	8-K of 20 Oct; ACCC register	◆ ◆
20 Oct 2026	All four foreign reviews reported complete, subject to the 14-day window for applications to the Australian Competition Tribunal	8-K, Item 8.01	◆
27 Oct 2026	Window closed (13 Oct plus 14 calendar days); no application on the Tribunal's list read on 10 Nov	Derived; Tribunal list	◆
29 Oct 2026	Stockholders adopt the agreement: 118,218,914 for, 7,592,263 against, 214,454 abstaining	8-K, Item 5.07	◆
30 Oct 2026	Continues to expect closing by the first quarter of 2027 (forecast); guidance withdrawn	Results release	◆
9 Nov 2026	Expects the parties to be in a position to close during the fourth quarter of 2026, subject to the remaining conditions (forecast)	8-K, Item 7.01	◆
10 Nov 2026	No 8-K under Item 2.01, Form 25, Form 15 or completion release: pending	EDGAR; investor feed	◆ ◆
6 May 2027	Outside date, extended to 6 Aug 2027 only if every non-regulatory condition is satisfied by then	8-K, Item 1.01	◆ [115]

The expected closing moved twice, from the end of 2026 to the first quarter of 2027 and back to the fourth quarter of 2026, and every statement of it is a forecast ◆. DoubleVerify withdrew its guidance and suspended earnings calls while the deal is pending ◆. On 10 November 2026, last checked at about 12:44 UTC, EDGAR held no closing 8-K, Form 25 or Form 15, and the investor feed no completion release, so DoubleVerify remained a listed company not owned by Nielsen's parent ◆. A closing before the freeze on 11 December 2026 will be dated from its own record. Stockholder complaints over the merger belong to Chapter 9 ◆.

Specialists sold, merged, renamed or retired

Protected by Mediaocean, formerly Protected Media, sits inside an ad-serving group: Flashtalking bought it in 2021, Mediaocean agreed that July to buy Flashtalking, and no completion date for that purchase, or date for the rename, was found ◆. Mediaocean announced in May 2022 that verification was fully integrated into Flashtalking's ad serving, and in February 2025 it completed its purchase of another ad server, Innovid ◆ ◆. Innovid's merger proxy of January 2025 calls Mediaocean a portfolio company of funds managed by CVC Capital Partners and of TA Associates and Charlesbank Capital Partners, wording that does not state control ◆.

HUMAN, formerly White Ops, renamed itself on 30 March 2021, three months after Goldman Sachs Merchant Banking Division, ClearSky Security and NightDragon bought it ◆. It then raised two rounds led by WestCap and announced a merger with PerimeterX on 27 July 2022, with approvals received and no completion date stated ◆. In a court filing docketed on 26 June 2026 HUMAN named Singularity Buyer LLC and Singularity TopCo LLC as its corporate parents; no source opened says who controls them ◆.

Adloox passed from Scope3 to Peer39 in June 2026, after Scope3's November 2024 acquisition ◆ ◆. Adloox Limited's September 2026 accounts name Peer39 Tech LLC as ultimate controlling party from June ◆. MRC's 29 June letter continued the renamed service's post-serve accreditation, with the full metric and environment scope in Chapter 2 ◆ ◆. Scope3's later rename to Apostra leaves Adloox's ownership unchanged ◆. Peer39 describes itself as independent and employee owned, but the sources screened to 10 November do not settle its group ownership; the UK subsidiary record covers that entity only ◆.

Oracle retired Moat instead of selling it: its end-of-life notice of 12 August 2024 ended every Oracle Advertising product and its support, Moat Analytics and Moat Reach among them, from 30 September 2024 ◆. No sale was found in the notice or in the one search logged, and a retirement does not show that Moat's measurement failed ◆. Pixalate, Zefr and Channel Factory recorded smaller events, listed in the table; Pixalate told a court in April 2025 that it has no corporate parent, and none of the three states who controls it ◆ ◆.

H5 is contradicted, narrowly

H5 fails on one page, and only narrowly. It has two parts. The first says both public verifiers' last filings show seller-side revenue, and it holds ◆ ◆. The second says no scored product with seller-side revenue, or with an owner that has an interest in the inventory or the currency it measures, publishes a conflict or separation policy covering that interest. The registered rule calls H5 contradicted if "at least one such policy is published by a scored product or its owner", and supported if "no published conflict or separation policy covering such an interest is found in the sources screened".

On 13 June 2024 DoubleVerify published a page on measuring publishers impartially. It says its media quality assessment teams work independently of its business development and client relations teams, calls that a structural separation that keeps its evaluations free of commercial influence, says it does not alter its methods or soften its assessments because of a publisher's status or relationship with DoubleVerify, and lets publishers raise site-classification concerns directly ◆. Read as the rule is written, that is a published separation statement covering DoubleVerify's seller-side interest in its publisher clients, and the paper codes H5 as contradicted on it ◆.

The separation page names no responsible policy owner or enforcement mechanism and does not address Neptune's prospective interests ◆. Its general audit assertion identifies no report tied to the separation ◆. DoubleVerify's Code of Ethics concerns personal conflicts and its 10-K discusses Providence's corporate opportunities; neither supplies the missing governance record ◆. Its Impact Report's claim of no inherent conflict is vendor positioning alongside seller revenue ◆. Requiring a policy with scope and enforcement would instead support H5 within the products screened ◆.

The merger release promises independent operating structures without a mechanism or duration, while also describing an integrated platform ◆. The definitive-proxy search found no conflict, firewall or separation policy ◆. The declared Samba–Nielsen competition interest applies here. This records governance disclosure and establishes no bias in DoubleVerify's measurements.

For the other vendors screened, no conflict or separation policy was found ◆ ◆ ◆. The table sets out what each discloses.

VENDOR	SELLER-SIDE OR PLATFORM CUSTOMERS	SELLER-SIDE AMOUNT	OWNER AND ITS INTERESTS	SAFEGUARD FOUND	EVIDENCE
DoubleVerify	Supply-side customer type: publishers, retail media networks, ad networks, exchanges	$71.3m of $748.3m, FY2025 (9.5%)	Listed; sale to Neptune BidCo US Inc., parent of the Nielsen Company (US), LLC, pending on 10 Nov 2026	Separation page of 13 Jun 2024 (publisher clients; no scope, owner or enforcement); nothing on the pending owner	◆ ◆ ◆
IAS	Publisher channel, including Publica's CTV ad serving; media partners can pay advertisers' fees	$76.5m of $530.1m, FY2024 (14.4%); FY2025 not filed	Novacap's funds since 23 Dec 2025; the portfolio includes Static Media (digital media) and Cadent (TV advertising)	None found	◆ ◆ ◆
Protected by Mediaocean	Supply-side platforms (SSPs), DSPs, exchanges and publishers among its customers in 2021	Not filed	Mediaocean, which also runs the Flashtalking and Innovid ad servers	None found	[116] ◆ ◆
Adloox	Publishers (ad-server segments; Brut since 2023)	Not filed	Peer39, which sells pre-bid controls and publisher categorisation	None found	◆ ◆
HUMAN	DSPs, SSPs and exchanges (platform suite announced 5 Oct 2026); Magnite since 2015	Not filed	Private; parents Singularity Buyer LLC and Singularity TopCo LLC; control not stated	None found	◆ ◆ ◆
Pixalate	SSPs and publishers among its customers	Not filed	Private; no corporate parent (April 2025)	Method rules only	◆ ◆
Zefr	Not established; an IRIS.TV revenue share reported in 2020	Not filed	Private; control not stated	None found	◆ ◆
Channel Factory (ViewIQ)	Not established; relationships with creators and channel owners	Not filed	Truelink Capital investment (2025); Star Mountain a minority owner	None found	◆ ◆

IAS's supply-side business includes publisher tools, Publica, Quality Connect, pre-bid segment resale and Media Partner Fees ◆. Novacap also lists Static Media and Cadent, but whether IAS measures their inventory is undocumented ◆. Searches of IAS's site and filings on 10 November found no policy covering these interests ◆.

Mediaocean owns two ad servers and the verifier whose data it integrates into ad serving, and frames its independence as freedom from media ownership ◆. Peer39 sells pre-bid controls and says its post-buy reporting exists to confirm that those controls worked, with no separate fee for it ◆. Zefr buys YouTube media for advertisers and in 2023 supplied video-level signals to TikTok's first-party inventory filter ◆. Channel Factory builds suitability lists and is paid to buy media against them, an inference labelled as such in the evidence ◆. Pixalate's published safeguards are method rules: its Seller Trust Index ranks on buy-side data only, its terms bar using a client's first-party data to rank that client, and its clawback labels cannot be changed by hand; none addresses Pixalate measuring a seller client's own traffic ◆. For the private vendors the amounts are unresolved under the rule, since no disclosure means unknown ◆ ◆.

These absences rest on the pages, filings and searches logged on 9 and 10 November 2026; vendors were not contacted, and a policy kept inside order forms or customer

documentation would not be visible ◆ ◆. The screen also missed some registered candidates. Among those not screened for a safeguard are Proximic by Comscore, illumaSentinel, mFilterIt, Fraudlogix, which sells to bidding platforms, and Escalated, so for them the absence is not established ◆ ◆ ◆ ◆ ◆. That gap cannot reverse the verdict as coded, which one published policy settles; under the strict reading it would limit any support to the products screened.

Interest disclosure is scored in Chapter 11. Ownership itself earns or loses no points ◆ ◆ ◆.

CHAPTER 8

After GARM

The brand-safety floor and suitability tiers of GARM, the Global Alliance for Responsible Media, outlived it. The World Federation of Advertisers (WFA) announced on 9 August 2024 that it was discontinuing GARM, the framework's last edition is dated June 2022, and the documents that still use a floor, from the Trustworthy Accountability Group's (TAG) certification guidelines to vendors' and Google's pages, point back to that edition, to undated copies of it or to older frameworks. One screened body publishes a later framework with both a floor and tiers: ISBA, the UK advertisers' trade body, whose Responsible Media Guide of March 2025 keeps ten of GARM's twelve floor categories under three risk tiers without naming GARM. On the rule's wording that contradicts H6, the hypothesis that no maintained successor exists, and the coding rests on one reading: that ISBA's pledge of regular review, which comes with no published change log or procedure, is the "stated change process" the rule asks for.

IN SHORT

The WFA discontinued GARM in August 2024 and said in July 2026 that it would not restart it ◆ ◆. Its framework page's last update is June 2022; BSI's hosted copy has twelve floor categories and risk tiers but no date or attribution ◆ ◆. TAG maintains a certification floor without tiers, while IAB Tech Lab maintains category labels ◆ ◆. ISBA's March 2025 brand safety chapter publishes ten floor categories and high, medium and low tiers; its hub names an overseer and pledges review, without a public procedure ◆ ◆. The author's reading of that pledge as a stated change process makes H6 contradicted ◆. Requiring a documented procedure makes it unresolved because the May 2026 full guide is member-only ◆. Vendor references to GARM do not establish which dated edition their products implement ◆ ◆ ◆.

GARM closed, and its framework stopped in 2022

The WFA announced GARM's discontinuation on 9 August 2024, citing the drain from allegations; its framework page gives 8 August ◆ ◆. X and Rumble had sued on 6 August, and Chapter 9 dates those cases ◆ ◆. The WFA's July 2026 promise not to restart GARM appeared in its settlement announcement, whose terms were unpublished ◆.

The WFA page records June 2022 as the framework's latest update, with the document members-only ◆. That edition has twelve floor categories and three risk tiers ◆. The BSI

copy used for textual comparison carries no date, version or attribution, so its identification with that edition is the paper's inference ◆ ◆.

On 15 April 2026 the Federal Trade Commission (FTC) and eight states alleged that Dentsu, WPP Media and Publicis used the APB and GARM to set a common "Brand Safety Floor" aimed at "misinformation"; the agencies settled the same day by consent orders, no verification vendor is a defendant, and the case is dated in Chapter 9 ◆.

Who publishes a floor now

Of the ten bodies H6 names (the WFA, the Association of National Advertisers (ANA), the 4A's APB, the IAB, IAB Tech Lab, TAG, the Media Rating Council (MRC), ISBA, BSI and the European Advertising Standards Alliance (EASA)), only ISBA publishes a floor and tiers under a date later than June 2022 in the documents screened to 10 November 2026 ◆. The others certify against the floor, keep categories without tiers, host undated copies, publish principles, or show no framework of their own dated after June 2022 ◆ ◆ ◆.

BODY	DOCUMENT READ	DATED AFTER JUNE 2022	FLOOR	TIERS	CHANGE LOG OR PROCESS	EVIDENCE
WFA (GARM)	Brand Safety Floor + Suitability Framework page; document members-only	No: "Last updated: June 2022"	12 categories	High, medium, low	None; GARM discontinued; no restart (29 July 2026)	◆
ISBA	Responsible Media Guide, brand safety chapter	V1, March 2025; guide "Updated May 2026" (member-only)	10 categories	High, medium, low	No log; pledge of regular review and update	◆ ◆
TAG	Brand Safety Certified Guidelines v5.0, Appendix B	July 2026	12 categories, plus 11 IAB Tech Lab Sensitive Topics at "Floor"	None	Change log	◆ ◆
4A's APB	Brand Safety Floor and Brand Suitability Framework	No: 21 September 2018	13 categories	Risk levels (2018)	None found; the 4A's link leads to an unrelated members-only page	◆ [117]
4A's, with the APB	Media Responsibility Principles, updated	11 December 2023	None (eleven principles)	None	Update announced; no log	◆ [118]
BSI	Hosted floor-and-tiers table	Undated	12 categories	High, medium, low	None; no owner or version	◆ [119]
IAB Tech Lab	Content Taxonomy 3.1; 3.0 Descriptive Vectors	3.1 posted 9 December 2024, last revised 22 July 2025; tier labels June 2022	11 Sensitive Topics	Labels only, undefined	Versioning with public comment	◆
IAB	Brand Safety and Suitability Guide	No: December 2020	As GARM and the APB then	As GARM and the APB then	None found	◆ [120]
ANA	Q3 2025 benchmark release	No framework in the release	Not published	"Low-risk" share only	Not applicable	◆ ◆
MRC	2018 content-level supplement	No framework	Refers to the APB's floor	None of its own	Not applicable	◆

BODY	DOCUMENT READ	DATED AFTER JUNE 2022	FLOOR	TIERS	CHANGE LOG OR PROCESS	EVIDENCE
EASA	Publication pages 1 to 5	No framework; a December 2023 repository describes GARM	Not applicable	Not applicable	Not applicable	[121]
UK Government Communication Service (outside the list)	SAFER framework for digital brand safety	2 March 2026, with a summary of changes	None (five principles)	None	Summary of changes	[122]
AgenticAdvertising.org (outside the list)	Common Sense Brand Standards 1.0.0	Effective 1 January 2026	Exclusion floor and graduated categories	None defined	No revisions recorded	

TAG keeps the floor inside a certification standard. Appendix B of its Brand Safety Certified Guidelines v5.0 lists the twelve "Brand Safety Floor" categories, linking to BSI's copy, and the eleven Sensitive Topics of IAB Tech Lab's Content Taxonomy version 2.2 at the "Floor" risk-rating, with no tiers and no mention of GARM [21]. Version 2.2 dates from December 2020; IAB Tech Lab posted version 3.1 for comment in December 2024 [123]. The guidelines keep a change log back to version 1.0 of September 2020 and certify companies, not products. H6 as registered counts a certification standard that reproduces the floor as certification maintenance unless it also publishes the tiers, and TAG's does not. In TAG's registry on 10 November 2026, DoubleVerify, IAS, Adloox, Channel Factory and illuma held the Brand Safety Certified seal, and Zefr was enrolled with no seal.

TAG links to BSI's undated, unattributed floor-and-tiers table. BSI also hosts the APB's 2018 framework and the 4A's July 2020 playbook [119, 124].

The accreditor's standard points to a framework whose publisher's own link now leads elsewhere. MRC's 2018 supplement on content-level brand safety, the standard behind its content-level accreditation, sets no suitability tiers of its own: it asks for reporting discrete enough for buyers to set their own risk tolerance and refers to the APB's floor and Brand Suitability Framework. The APB released that framework on 21 September 2018, with thirteen categories [117]. On 10 November 2026 the 4A's link labelled "Brand Safety Floor and Suitability Framework" led to an unrelated members-only page, and the only dated APB-related output found since June 2022 was the Media Responsibility Principles of 11 December 2023, updated by the 4A's Media, Tech and Data practice with the APB: eleven principles, with no floor categories or tiers.

The latest brand safety and suitability guide found on iab.com, from the IAB, the trade body of digital media sellers and platforms, is dated December 2020 [120]. IAB Tech Lab maintains categories without tiers: its Content Taxonomy 3.1, posted for public comment on 9 December 2024 and last revised on 22 July 2025, keeps the eleven Sensitive Topics derived from GARM under versioning with public comment but has no tier rows, and the tier labels Floor, High, Medium and Low Risk appear, undefined, in older files: version 2.2 of December 2020 and the 3.0 Descriptive Vectors of June 2022 [123].

The ANA reports suitability results without publishing, in the releases read, the framework behind them. Its programmatic benchmark, built with TAG TrustNet and Fiducia on log-level data, added brand safety and suitability metrics in Q3 2025 from what it calls "an AI-powered standard" and reported 99.1% of the benchmark's programmatic spend in low-risk environments, without saying who classifies "low-risk" or against which framework or tiers; the benchmark report itself was not opened.

Two dated documents outside H6's list stop short of a floor-and-tiers framework. The UK Government Communication Service's SAFER framework of 2 March 2026 updates its 2020 SAFE framework with a summary of changes, but it is one advertiser's five principles, with no floor or tiers [122] ◆. AgenticAdvertising.org's policy registry serves Common Sense Brand Standards 1.0.0, a standard that originated at Scope3 (now Apostra) and took effect on 1 January 2026, with an exclusion floor and "graduated assessment" categories but no defined tiers; its history showed no revisions on 10 November 2026 ◆ ◆. The author co-leads that organisation's Signals & Measurement working group, a declared interest; agent buying is the subject of Chapter 10 ◆.

JICWEBS went into TAG

The UK's Joint Industry Committee for Web Standards (JICWEBS), which once ran UK digital trading standards, now sits inside TAG, and the one floor-and-tiers framework found with a date after June 2022 comes from ISBA, one of its four former controllers ◆ ◆. JICWEBS was controlled by the UK trade bodies of online publishers (AOP), digital advertising (IAB UK), agencies (IPA) and advertisers (ISBA); its Digital Trading Standards Group (DTSG), set up in 2012, asked buyers and sellers to state in contract where ads should or should not appear, using independently certified content verification tools or schedules (lists) ◆.

DATE	EVENT	EVIDENCE
2012	JICWEBS's Digital Trading Standards Group set up	◆
31 August 2016	JICWEBS Ltd incorporated, controlled by AOP, IAB UK, IPA and ISBA	◆ ◆
March 2017	JICWEBS and TAG announce a formal partnership	◆
January 2019	TAG's Certified Against Fraud replaces JICWEBS's anti-fraud principles as the UK's main anti-fraud standard	◆
21 January 2020	Content Verification Guide issued by AOP, IAB UK, IPA, ISBA, JICWEBS and Newsworks	◆
29 July 2020	TAG and JICWEBS announce their merger, for completion in early September	◆
1 September 2020	TAG recorded at Companies House as JICWEBS Ltd's controller (75% or more of votes); the four trade bodies cease	◆
10 September 2020	TAG Brand Safety Certified launches, version 1.0 merging the DTSG with TAG's Inventory Quality Guidelines; 112 companies certified at launch	◆ ◆
March 2025	ISBA's Responsible Media Guide, brand safety chapter V1	◆
May 2026	ISBA's page marks the guide "Updated May 2026" (member-only)	◆
July 2026	TAG Brand Safety Certified Guidelines v5.0	◆
10 November 2026	JICWEBS Ltd still "Active" at Companies House, its confirmation statement overdue; jicwebs.org redirects to tagtoday.net	◆

The merger also drew the UK bodies toward TAG's governance: the 2020 announcement said AOP, IAB UK, IPA and ISBA would join TAG's Leadership Council, and ISBA's guide says ISBA sits on that council and on MRC's board of directors, which this paper records as an interest in the bodies that grade verifiers, not as evidence of bias ◆. ISBA's guide advises members rather than setting a trading standard: it calls an ad verification partner "the principal way to manage brand safety", names Adloox, DoubleVerify, IAS, Mantis and Zefr as examples, and says its guidance often feeds ISBA's member-only Media Services Framework for agency contracts ◆.

ISBA's guide against the rule

ISBA's guide meets the rule's date, floor and tier elements on ISBA's own documents; only the change-process element needs a reading ◆. The brand safety chapter, produced with The Aperto Partnership, an advisory firm to advertisers, carries "V1 March 2025" on its cover and calls the guidance "entirely voluntary" ◆ [10]. It sets a floor of ten content categories with descriptions, a suitability framework of high, medium and low risk "against the same ten categories", and an adjacency framework ◆. Its descriptions match the unattributed, undated floor-and-tiers text BSI hosts, which this paper identifies as GARM's June 2022 framework; the WFA's own copy is members-only and was not compared ◆. The guide does not name GARM, drops two of GARM's twelve categories (debated sensitive social issues and misinformation), presents the floor as "industry-accepted" and ties the tiers to the inventory filters platforms already offer ◆.

ELEMENT OF THE RULE	WHAT ISBA'S DOCUMENTS SHOW	READING	EVIDENCE
Published by a screened body	ISBA is one of the ten bodies H6 names	Met	◆
A dated edition later than June 2022	Chapter cover "V1 March 2025"; hub page dated 20 March 2025 and marked "Updated May 2026"	Met	◆
Floor categories	Ten, with descriptions	Met	◆
Suitability tiers	High, medium and low against the same ten	Met	◆
A change log or a stated change process	No change log, decision procedure or review cycle published; a pledge that the guide will be "regularly reviewed and updated", oversight by ISBA's Media Leaders group, and a promise to update after further Ofcom guidance	The judgement	◆ ◆

The remaining judgment concerns what the rule means by a change process. ISBA names oversight by its Media Leaders group, pledges regular review and identifies further Ofcom guidance as an update trigger ◆ ◆. Its hub records a May 2026 update to the member-only full guide, whose effect on the public March 2025 floor and tiers was not readable ◆.

The author treats the pledge, overseer and trigger as the rule's stated change process and codes H6 contradicted ◆. A requirement for decision rights and a documented procedure instead yields unresolved ◆. ISBA is a screened body under the rule, although its voluntary UK guide leaves the hypothesis statement's broader cross-industry wording open to challenge ◆.

What GARM alignment means now

Where a scored product names a framework on the pages read, it names GARM, an IAB category list or an unnamed industry standard, and none names an edition ◆ ◆ ◆ ◆. A House Judiciary Committee staff report of July 2024, a characterisation rather than a finding, said committee documents suggested GARM had been building its framework into the products of the ad-tech providers it invited, among them DoubleVerify, IAS, Peer39, Channel Factory, Pixability and Zefr (Chapter 9) ◆.

PRODUCT (COHORT)	FRAMEWORK NAMED	EDITION NAMED	EVIDENCE
IAS advertiser verification (A)	Low, Medium or High "across the supported 12 GARM content categories"; social-platform models "trained against the GARM Framework"; a separate open-web taxonomy of 8 categories and 5 risk thresholds	None	◆ ◆
DoubleVerify advertiser verification (A)	Its own Brand Safety Floor and Suitability Tiers, "mapped to the GARM framework" (January 2024 release); over 100 avoidance categories; no product page read names GARM	None	◆
Adloox verification (A; owned by Peer39 since June 2026)	"Based on the GARM framework" in Google's DV360 help; 12 exclusion segments on Adloox's DV360 page; "industry-standard BS&S framework" in its Meta release	None	◆ ◆
mFilterIt Ad Traffic Validation (A)	Brand-safety categories "as per the GARM guidelines"	None	◆
HUMAN Ad Integrity Suite (B)	None on its brand-safety page (captured 4 February 2026)	None	◆
Zefr brand safety and suitability (C)	GARM by name through 2023; twelve "industry-aligned" categories from September 2024; "GARM-aligned" on its Media Services page	None	◆
Peer39 brand safety and suitability (C)	"The industry standard brand safety and suitability framework, formerly known as GARM" (30 October 2026)	None	◆
illumaSentinel (C)	None named on illuma's pages read on 10 November 2026	None	[125]
ViewIQ, Channel Factory (C)	"IAB standard category lists"; "GARM+ strategies", undefined	None	◆

IAS's brand safety and suitability overview, published on 6 March 2024 and still online when read on 10 November 2026, scores suitability "across the supported 12 GARM content categories" and says its social-platform models are "trained against the GARM Framework"; its help-centre articles on GARM sit behind a login and were not read ◆ ◆ [126]. Google's pages carry the name too. DV360's help, read on 10 November 2026, still offers Adloox's pre-bid filters under "Scope3 (formerly Adloox)", "based on the GARM framework", although Peer39 has owned Adloox since June 2026 ◆ ◆. Google's Advertising Technology Partners directory describes its brand-safety partners' verification as "based on global industry standards defined by" GARM ◆ [127]. Peer39, writing on 30 October 2026, calls its basis "the industry standard brand safety and suitability framework, formerly known as GARM", and names no maintainer ◆.

Other vendors have dropped the name and kept the count. Zefr's releases mapped content to eleven GARM categories at its TikTok launch in January 2022 and to every GARM category, misinformation included, from February 2023; its September 2024 YouTube and October 2024 Snap releases describe twelve "industry-aligned" categories without naming GARM ◆. DoubleVerify called its own floor and tiers "mapped to the GARM framework" in January 2024, and none of its product pages read on 10 November 2026 names GARM or publishes category definitions ◆. Channel Factory's "GARM+" cannot be mapped to the June 2022 edition from its pages ◆.

The WFA page documents no official GARM edition later than June 2022 ◆. Counts and labels do not show which definitions a vendor implements or whether it has adapted them ◆ ◆. The public floors also differ, with ten categories at ISBA, twelve at TAG and BSI, eleven IAB Tech Lab Sensitive Topics, and thirteen in the APB's 2018 floor. A framework name alone cannot identify the categories used ◆ ◆ ◆.

H6 under its rule

The primary H6 verdict remains contradicted at medium confidence on the author's interpretation of ISBA's review pledge ◆. The ten-body screen and the table above document the date, floor and tiers; treating a process as requiring published decision rights instead yields unresolved ◆ ◆.

The other nine screened bodies supplied no qualifying edition in the documents read to 10 November 2026 ◆ ◆ ◆ ◆. H6's separate implementation clause cannot be settled from the product pages, which name no dated edition; the apparent June 2022 mapping remains an inference and GARM+ is unmapped ◆ ◆ ◆. Both the framework record and the primary judgment are due for recheck on 11 December ◆ ◆.

CHAPTER 9

The record, 2023 to 2026

From January 2023 to 10 November 2026 the public record on the scored products is letters, case reports, complaints, pleading-stage rulings, voluntary dismissals, consent orders against advertising agencies and the parties' own statements. In the sources screened, no court, regulator or accreditor made a finding on the merits about any scored product's accuracy; the only regulator decision on the merits in the registers searched, Australia's clearance of DoubleVerify's sale, concerns competition. Of the matters still open, only DoubleVerify's suit against Adalytics turns on how a verification product performs, and its one ruling took DoubleVerify's allegations as true. The record sits outside the Verdict Quotient (VQ), the paper's score: nothing in it scores. Its absences are bounded to the public sources searched on 10 November 2026, beyond which lie nonpublic inquiries, sealed filings and state dockets.

IN SHORT

Senators sent eight letters in two rounds before DoubleVerify sued Adalytics in May 2025; no action on the requests was found in the screened agency sources ◆ ◆ ◆. The April 2026 ruling in that suit accepted the complaint's allegations for the motion and made no factual finding about the report's accuracy ◆. DoubleVerify's securities class action was voluntarily dismissed without prejudice; IAS's was dismissed with prejudice and appealed ◆ ◆. The FTC withdrew its investigative demand to a third party in April 2026 and settled with three agency groups, with no verifier named as a defendant ◆ ◆. X's and Rumble's GARM cases ended in pleading-stage dismissals; after the WFA settled, X's appeal continued against other defendants ◆ ◆. Australia's clearance of DoubleVerify's proposed sale decides competition and establishes no measurement accuracy ◆.

What the record is

Q1 records each letter, complaint, ruling and response by its legal posture and primary date, outside the VQ. Its fixed rule keeps allegations, pleading decisions and voluntary dismissals apart from factual findings. A request for an investigation does not show that one opened. Unread documents are identified, including Kaszirer's dismissal term and the Delaware IAS-share case ◆ [128]. The window starts on 1 January 2023; pending matters were checked on 10 November 2026 and are due for recheck on 11 December ◆ ◆.

The register

The table lists, oldest first, each record found that concerns a scored product or GARM, the WFA's Global Alliance for Responsible Media.

DATE	FORUM	DOCUMENT	WHAT IT IS	STATUS AT 10 NOVEMBER 2026	EVIDENCE
4 May 2023	Autorité de la concurrence (France)	Decision 23-MC-01 on Adloox's request for interim measures against Meta	Ruling (interim measures; merits not decided)	Final, not appealed; no merits decision listed	◆
21 September 2023	C.D. Cal. 2:23-cv-07914	Montgomery Venkersammy v. Human Security, removed to federal court	Complaint (subject not established)	Reported settled by mediation; terminated 22 May 2024	◆
5 February 2024	D. Minn. 0:24-cv-00310	Thor v. DoubleVerify	Complaint (subject not established)	Stayed 23 February 2024; dismissed 27 January 2025	[129]
18 March 2024 (public by)	Adalytics	Report on ads on made-for-advertising sites in January 2024	Independent report (cases)	Published	◆
4 April 2024 (public by)	Adalytics	Report on the Forbes www3 subdomain	Independent report (one case)	Published	◆
10 May 2024 (first capture)	Adalytics	Report on user-ID rotation at Colossus, naming HUMAN only as the supply-side platform's declared partner	Independent report (one case)	Published	◆
10 July 2024	House Judiciary Committee, majority staff	"GARM's Harm", interim staff report	Congressional staff report	Final at publication	◆
6 August 2024	N.D. Tex. 7:24-cv-00114	X Corp. v. World Federation of Advertisers et al.	Complaint	Dismissed 26 March 2026	◆
6 August 2024	N.D. Tex. 7:24-cv-00115	Rumble v. World Federation of Advertisers et al.	Complaint	Dismissed 13 August 2025	◆
7 August 2024	Adalytics	"AI" brand safety technology	Independent report (cases)	Published	◆
9 August 2024	WFA	Statement discontinuing GARM	Party statement	GARM discontinued	◆ ◆
11 October 2024	Marketing Brew	Report of questions from Justice Department and Naval Criminal Investigative Service officials to advertising executives about IAS and DoubleVerify	Press report of a nonpublic inquiry	Unverified; no primary record	◆
29 January 2025	S.D.N.Y. 1:25-cv-00847	Oklahoma Firefighters Pension and Retirement System v. Integral Ad Science	Securities class action complaint	Dismissed 2 March 2026; on appeal	◆
7 February 2025	Senators Blackburn and Blumenthal	Letter to DoubleVerify	Letter (five questions)	Answers asked by 14 February 2025; none found in the pages opened	◆
7 February 2025	Senators Blackburn and Blumenthal	Letter to IAS	Letter (questions)	IAS says it answered; the answer was not found	◆ [79]
7 February 2025	Senators Blackburn and Blumenthal	Letter to Amazon	Letter (dated from the release; text not read)	Answer not searched for	[130]

DATE	FORUM	DOCUMENT	WHAT IT IS	STATUS AT 10 NOVEMBER 2026	EVIDENCE
7 February 2025	Senators Blackburn and Blumenthal	Letter to Google	Letter (dated from the release; text not read)	Answer not searched for	[130]
7 February 2025	Senators Blackburn and Blumenthal	Letter to the Media Rating Council	Letter (dated from the release; text not read)	Answer not searched for	[130]
7 February 2025	Senators Blackburn and Blumenthal	Letter to the Trustworthy Accountability Group	Letter (dated from the release; text not read)	Answer not searched for	[130]
7 February 2025	Adalytics	Report on ads on an image host (imgbb.com, ibb.co)	Independent report (cases; advertisers' statements relayed)	Published	◆
11 March 2025	S.D.N.Y. 1:25-cv-02039	Tauber derivative complaint against IAS directors, consolidated with Newman (filed 2 June 2025)	Derivative complaints	Voluntarily dismissed 15 January 2026	◆
24 March 2025	Senator Warner	Letter to the FTC chairman, released 28 March	Letter (request to investigate; answers by 28 April 2025)	No FTC action found	◆
24 March 2025	Senator Warner	Letter to the Attorney General, released 28 March	Letter (request to investigate; no deadline)	No Justice Department action found; its releases not searched separately	◆
28 March 2025	Adalytics	Report on pre-bid bot detection and filtration	Independent report (cases)	Subject of DoubleVerify v. Adalytics	◆
31 March 2025	S.D.N.Y. 1:25-cv-02715	Martin v. Zagorski, pro se, against DoubleVerify, Providence Equity and four individuals	Complaint (whistleblower retaliation, civil rights)	Dismissed 16 September 2026	◆
30 April 2025	S.D.N.Y. 1:25-cv-03614	Pixalate v. Criteo, Pixalate as plaintiff	Contract complaint (subject not established)	Voluntarily dismissed 18 June 2025	◆
12 May 2025	D. Md. 8:25-cv-01535	DoubleVerify v. Adalytics	Complaint (Lanham Act, defamation, injurious falsehood, tortious interference, unfair competition)	Partly dismissed 27 April 2026; continuing	◆
20 May 2025	FTC, File No. 251-0061	Civil investigative demand to the Global Disinformation Index (GDI), listing IAS and DoubleVerify among thirteen rating entities	Investigative demand to a third party	Withdrawn 16 April 2026	◆
22 May 2025	S.D.N.Y. 1:25-cv-04332	Securities class action against DoubleVerify	Securities class action complaint	Voluntarily dismissed 22 December 2025	◆
13 August 2025	N.D. Tex. 7:24-cv-00115	Rumble v. WFA, opinion	Ruling (personal jurisdiction and venue; merits not reached)	Closed; no appeal or refiling found at the 9 November check	◆
9 December 2025	S.D.N.Y. 1:25-cv-10200	Kaszirer v. Zagorski, against ten DoubleVerify officers and directors	Derivative complaint	Notice of voluntary dismissal filed 4 February 2026, entered 6 February; prejudice term not read	◆
22 December 2025	S.D.N.Y. 1:25-cv-04332	Lead plaintiffs' notice in the DoubleVerify class action	Voluntary dismissal without prejudice	Closed 23 December 2025	◆
3 January 2026	arXiv (New Jersey Institute of Technology; DeepSee.io)	Smith et al., consistency of brand-safety classifications	Independent study (agreement between vendors)	Published	◆
15 January 2026	S.D.N.Y. 1:25-cv-02039	Stipulation and order in the IAS derivative litigation	Voluntary dismissal (terms not read)	Closed	◆

DATE	FORUM	DOCUMENT	WHAT IT IS	STATUS AT 10 NOVEMBER 2026	EVIDENCE
2 February 2026	D.D.C. 1:25-cv-04137	GDI's amended complaint against the FTC, alleging IAS received a demand	Complaint (pleaded on information and belief)	Closed 4 May 2026	◆
February 2026	C.D. Cal. 2:26-cv-01108	Singer v. Comscore, privacy class action (not shown to concern Proximic)	Complaint	Dismissed with leave to amend 30 June 2026	◆
26 February 2026	Delaware Court of Chancery, per Law360	A Glazer Capital fund's suit for payment for its IAS shares	Press report of a complaint	Unverified	[128]
2 March 2026	S.D.N.Y. 1:25-cv-00847	Oklahoma Firefighters v. IAS, opinion	Ruling (pleading stage; dismissal with prejudice)	Judgment 5 March; appeal noticed 26 March; briefing per the plaintiffs' firm (undated)	◆
26 March 2026	N.D. Tex. 7:24-cv-00114	X Corp. v. WFA, opinion	Ruling (no antitrust injury, with prejudice for the WFA and domestic defendants; personal jurisdiction, without prejudice, for four foreign defendants)	On appeal, 5th Cir. 26-10394	◆
15 April 2026	FTC and eight states; N.D. Tex. 4:26-cv-00469	FTC v. Dentsu, WPP Media and Publicis	Complaint; consent orders (settlement, no admission)	Orders entered; Dentsu's runs ten years	◆
16 April 2026	FTC	Letter withdrawing GDI's demand; notice that several other demands were withdrawn	Withdrawal	GDI's suit closed 4 May 2026	◆
27 April 2026	D. Md. 8:25-cv-01535	DoubleVerify v. Adalytics, opinion	Ruling (pleading stage; tortious interference dismissed without prejudice)	Answer filed; status report due 9 February 2027	◆
5 May 2026	S.D.N.Y. 1:26-cv-03727	Chavez v. Human Security, privacy class action	Complaint	Voluntarily dismissed with prejudice, so ordered 21 July 2026	◆
11 May 2026	D. Md. 8:25-cv-01535	Adalytics' answer	Answer (pleads truth; calls the suit a strategic lawsuit against public participation)	Case continues	◆
29 July 2026	WFA	Release on the settlement with X	Settlement (terms not public); party statement	X's appeal continues against the other defendants	◆
5 August 2026	5th Cir. 26-10394	Stipulation dismissing the appeal and cross-appeal between X and the WFA (Doc. 119)	Stipulation (settlement)	Granted 11 August 2026, per the 9 November check	[131]
12 August 2026	5th Cir. 26-10394	United States amicus brief supporting X (Doc. 126)	Amicus brief	Briefing	◆
12 August 2026	5th Cir. 26-10394	West Virginia and 11 other states' amicus brief supporting reversal (Doc. 123)	Amicus brief	Briefing	◆
16 September 2026	S.D.N.Y. 1:25-cv-02715	Martin v. Zagorski, opinion	Ruling (dismissal with prejudice, failure to state a claim)	Judgment 17 September; no notice of appeal on the docket	◆
8 October 2026	N.Y. County Supreme Court	Kent v. DoubleVerify Holdings et al., seen only through an aggregator	Summons and complaint (not read)	Open per the aggregator; not confirmed as one of DoubleVerify's two complaints	◆
13 October 2026	ACCC, MN-55036	Phase 1 approval of Neptune BidCo's acquisition of DoubleVerify	Competition determination	Assessment completed	◆

DATE	FORUM	DOCUMENT	WHAT IT IS	STATUS AT 10 NOVEMBER 2026	EVIDENCE
20 October 2026	DoubleVerify, 8-K filing	Fourteen demand letters and two unserved stockholder complaints over the merger	Issuer filing reporting allegations	Unserved at 30 October 2026; DoubleVerify disputes them	◆

The dated record, 2023 to 2026

Every regulator action, congressional letter, court filing, issuer filing and independent report in the register, dated from its primary document and coded by what it is. Letters are not findings, and a ruling is procedural unless it says otherwise. Nothing here scores in the VQ, and nothing here is legal advice.

Machine-readable copy: data/rules.csv.

FORUM OR ISSUER	RECORD	HOW THE PAPER CODES IT, AND ITS POSTURE	DATE	SCORED PRODUCTS NAMED	DOCUMENT	SRC
C.D. Cal.	Montgomery Venkersammy v. Human Security, Inc. (2:23-cv-07914)	allegation: terminated 22 May 2024 (docket mirror)	2023-09-21	HUMAN	complaint (subject not established)	[135]
D. Minn.	Thor v. DoubleVerify Inc. (0:24-cv-00310)	allegation; procedural ruling (stay 23 Feb 2024; dismissal 27 Jan 2025): closed 27 Jan 2025	2024-02-05	DoubleVerify	complaint (subject not established)	[129]
Adalytics	Adalytics report on made-for-advertising sites, January 2024 (adalytics.io/blog/ads-observed-on-made-for-advertising-sites-in-january-2024)	allegation: published	2024-03-18	IAS; DoubleVerify; HUMAN; Pixalate; Oracle Moat	independent report (case observations)	[164]
Adalytics	Adalytics report on the Forbes www3 subdomain (adalytics.io/blog/ads-observed-on-www3-forbes-subdomain)	allegation: published; revised after first publication	2024-04-04	DoubleVerify; IAS; HUMAN; Oracle Moat (mentions)	independent report	[166]
Adalytics	Adalytics report on user-ID rotation (Colossus SSP) (adalytics.io/blog/user-id-rotation)	allegation: published	2024-05-10	HUMAN (as the SSP's declared partner)	independent report	[77]
D. Md.	Colossus Media, LLC v. Adalytics Research, LLC (8:24-cv-01402)	allegation: not established (CourtListener blocked)	2024-05-14	none (researcher sued by an SSP)	complaint (Lanham Act, defamation)	[135]
House Judiciary Committee	GARM's Harm (interim staff report) (staff report, 39 pp.)	allegation: final at publication	2024-07-10	DoubleVerify; IAS; Peer39; Channel Factory; Pixability; Zefr; Oracle Moat	congressional staff report	[283]
N.D. Tex.	X Corp. v. World Federation of Advertisers et al. (complaint) (7:24-cv-00114)	allegation: see 2026-03-26	2024-08-06	none (GARM)	complaint	[284]
Adalytics	Adalytics '"AI" brand safety technology' (adalytics.io/blog/ai-brand-safety)	allegation: published	2024-08-07	DoubleVerify; IAS	independent report	[167]
World Federation of Advertisers	WFA discontinues GARM (about-garm statement)	party statement: GARM discontinued	2024-08-09	none	party statement	[285]
Marketing Brew (anonymous sources)	Reported DOJ and NCIS questions to ad executives about IAS and DoubleVerify	unverified (outside reach): no primary record	2024-10-11	IAS; DoubleVerify	press report of a nonpublic inquiry	[286]
DoubleVerify	DoubleVerify pre-emptive statement on the forthcoming bot report (Setting the Record Straight)	party statement: live (modified 15 Oct 2025)	2025-01-10	DoubleVerify	vendor statement	[174]
S.D.N.Y.	Oklahoma Firefighters Pension and Retirement System v. IAS (complaint) (1:25-cv-00847)	allegation: see 2026-03-02	2025-01-29	IAS (company)	securities class action complaint	[287]

FORUM OR ISSUER	RECORD	HOW THE PAPER CODES IT, AND ITS POSTURE	DATE	SCORED PRODUCTS NAMED	DOCUMENT	SRC
Senators Blackburn and Blumenthal	Six Senate letters (Amazon, Google, DoubleVerify, IAS, MRC, TAG) (letters dated 7 Feb 2025)	allegation (questions): answers requested by 14 Feb 2025; none public found	2025-02-07	DoubleVerify; IAS	congressional letters	[130]
Adalytics	Adalytics image-host (imgbb.com, ibb.co) report (adalytics.io/blog/adtech-vendors-csam-full-report)	allegation: published	2025-02-07	DoubleVerify; IAS	independent report	[169]
IAS (SEC filing)	IAS FY2024 10-K, Item 3: no material proceeding (10-K FY2024)	party statement: superseded by 10-Q disclosure	2025-02-28	IAS	issuer filing	[27]
S.D.N.Y.	Tauber derivative complaint (IAS directors) (1:25-cv-02039)	allegation: see 2026-01-15	2025-03-11	IAS (company)	derivative complaint	[288]
Senator Mark R. Warner	Warner letters to the FTC and DOJ (two signed letters)	allegation (requests to investigate): no agency action found	2025-03-24	IAS; DoubleVerify; HUMAN	congressional letters	[132]
Adalytics	Adalytics pre-bid bot report (adalytics.io/blog/prebid-bot-filtration)	allegation: subject of DoubleVerify v. Adalytics	2025-03-28	DoubleVerify; IAS; HUMAN	independent report	[173]
S.D.N.Y.	Martin v. Zagorski (complaint) (1:25-cv-02715)	allegation: see 2026-09-16	2025-03-31	DoubleVerify (company and officers)	pro se complaint	[289]
S.D.N.Y.	Pixalate, Inc. v. Criteo S.A. (Pixalate as plaintiff) (1:25-cv-03614)	allegation: voluntarily dismissed 18 Jun 2025; terminated 20 Jun 2025	2025-04-30	Pixalate (plaintiff)	contract complaint	[290]
D. Md.	DoubleVerify v. Adalytics (complaint) (8:25-cv-01535)	allegation: see 2026-04-27	2025-05-12	DoubleVerify (plaintiff)	complaint (Lanham Act, defamation and others)	[291]
FTC (File No. 251-0061)	FTC civil investigative demand to GDI naming IAS and DoubleVerify among rating entities (CID specification 10)	regulator action (investigative step): withdrawn 16 Apr 2026	2025-05-20	IAS; DoubleVerify (named, not addressed)	regulator investigative demand to a third party	[292]
S.D.N.Y.	DoubleVerify securities class action (complaint) (1:25-cv-04332)	allegation: see 2025-12-22	2025-05-22	DoubleVerify (company)	securities class action complaint	[134]
S.D.N.Y.	Newman v. Utzschneider (derivative, IAS directors) (1:25-cv-04622)	allegation: consolidated 12 Jun 2025	2025-06-02	IAS (company)	derivative complaint	[135]
S.D.N.Y.	Amended class complaint against IAS adds Vista fund entities (1:25-cv-00847 ECF 30)	allegation: see 2026-03-02	2025-06-16	IAS (company)	amended complaint	[287]
N.D. Tex.	Rumble v. WFA dismissed without prejudice (jurisdiction and venue) (7:24-cv-0115 ECF 69)	procedural ruling: closed; no appeal found (9 Nov check)	2025-08-13	none (GARM)	court opinion	[293]
IAS (SEC filing)	IAS Q3 2025 10-Q discloses the class action (10-Q Q3 2025, Note 14)	party statement: IAS's last periodic report	2025-11-04	IAS	issuer filing	[294]
S.D.N.Y.	Kaszirer v. Zagorski (derivative complaint) (1:25-cv-10200)	allegation: see 2026-02-06	2025-12-09	DoubleVerify (company)	derivative complaint	[295]
IAS	IAS drops GDI input from Context Control Avoidance (news release)	party statement: live	2025-12-18	IAS	vendor statement (method change)	[296]
S.D.N.Y.	DoubleVerify class action voluntarily dismissed without prejudice (1:25-cv-04332 ECF 35)	voluntary dismissal (neither vindication nor liability): closed 23 Dec 2025	2025-12-22	DoubleVerify	notice of voluntary dismissal	[297]
NJIT; DeepSee.io	Smith et al., brand-safety classification consistency study (arXiv:2601.01303)	independent report (agreement estimand): published (ICWSM 2026 per the brief)	2026-01-03	DoubleVerify; IAS; Oracle	independent academic study	[83]
S.D.N.Y.	IAS derivative litigation dismissed by stipulation and order (1:25-cv-02039 ECF 19)	voluntary dismissal (terms not read): closed	2026-01-15	IAS	stipulation and order of voluntary dismissal	[298]

FORUM OR ISSUER	RECORD	HOW THE PAPER CODES IT, AND ITS POSTURE	DATE	SCORED PRODUCTS NAMED	DOCUMENT	SRC
D.D.C.	GDI amended complaint alleges IAS received an FTC CID (1:25-cv-04137 ECF 29)	allegation (information and belief): case closed 4 May 2026	2026-02-02	IAS	complaint against the FTC	[299]
S.D.N.Y.	Kaszirer notice of voluntary dismissal entered as signed (1:25-cv-10200 ECF 5)	voluntary dismissal (prejudice term not read): closed	2026-02-06	DoubleVerify	notice of voluntary dismissal	[300]
C.D. Cal.	Singer et al. v. Comscore, Inc. et al. (privacy class action) (2:26-cv-01108)	allegation: pending at 26 Mar 2026 (10-K)	2026-02	Comscore (company; Proximic not shown)	putative class action complaint	[301]
Law360 (secondary)	Glazer Capital fund sues for payment for IAS shares (Delaware Chancery)	allegation (unverified): not established	2026-02-26	IAS	press report of a complaint	[128]
S.D.N.Y.	Oklahoma Firefighters v. IAS dismissed with prejudice (1:25-cv-00847 ECF 62)	procedural ruling (pleading stage): judgment 5 Mar 2026; appeal noticed 26 Mar 2026	2026-03-02	IAS	court opinion	[302]
N.D. Tex.	X Corp. v. WFA dismissed (with prejudice for WFA and domestic defendants) (7:24-cv-00114 ECF 255)	procedural ruling (failure to state a claim): on appeal, 5th Cir. 26-10394	2026-03-26	none (GARM)	court opinion	[284]
S.D.N.Y. to 2d Cir.	Notice of appeal in Oklahoma Firefighters v. IAS (1:25-cv-00847 ECF 64)	procedural: being briefed (plaintiffs' firm, undated)	2026-03-26	IAS	notice of appeal	[287]
FTC; N.D. Tex.	FTC and eight states v. Dentsu, WPP Media and Publicis; consent orders (4:26-cv-00469)	allegation; settlement (consent order, no admission): orders entered; 10-year term	2026-04-15	none (agencies; brand-safety floor)	regulator complaint and consent orders	[303]
FTC	FTC withdraws GDI's CID; 'several other CIDs' withdrawn (ECF 37 and 37-1, D.D.C. 1:25-cv-04137)	regulator action (withdrawal): GDI case closed 4 May 2026	2026-04-16	none named	regulator letter and notice	[304]
D. Md.	DoubleVerify v. Adalytics: motion to dismiss granted only as to tortious interference (8:25-cv-01535 ECF 30)	procedural ruling (pleading stage): answer filed; status report due 9 Feb 2027	2026-04-27	DoubleVerify	court opinion	[305]
S.D.N.Y.	Chavez v. Human Security, Inc. (privacy class action) (1:26-cv-03727)	allegation: see 2026-07-21	2026-05-05	HUMAN	putative class action complaint	[306]
D. Md.	Adalytics answer (truth; anti-SLAPP) (8:25-cv-01535 ECF 32)	party statement: case continues	2026-05-11	DoubleVerify	answer	[307]
S.D.N.Y.	Chavez v. HUMAN voluntarily dismissed with prejudice (1:26-cv-03727 ECF 18)	voluntary dismissal (with prejudice; no finding): closed	2026-07-21	HUMAN	notice of voluntary dismissal, so ordered	[308]
World Federation of Advertisers	WFA and X settle; WFA will not restart GARM (release)	settlement (terms not public) and party statement: appeal continues against other defendants	2026-07-29	none	party statement	[309]
5th Cir.	X and WFA stipulate to dismissal of their appeal and cross-appeal (26-10394 Doc. 119)	settlement (terms not public): granted 11 Aug 2026 (9 Nov check)	2026-08-05	none	stipulation	[131]
5th Cir.	United States amicus supporting X (26-10394 Doc. 126)	party-side position: briefing	2026-08-12	none	amicus brief	[310]
5th Cir.	West Virginia and 11 states amicus supporting reversal (26-10394 Doc. 123)	party-side position: briefing	2026-08-12	none	amicus brief (state AGs)	[311]
D. Md.	DoubleVerify v. Adalytics: next status report due 9 Feb 2027 (8:25-cv-01535 ECF 40)	procedural ruling: pending	2026-09-02	DoubleVerify	scheduling order	[312]
S.D.N.Y.	Martin v. Zagorski dismissed with prejudice (1:25-cv-02715 ECF 61)	procedural ruling (pleading stage; judgment for defendants): judgment 17 Sep 2026; no appeal on docket at 10 Nov 2026	2026-09-16	DoubleVerify (company and officers)	court opinion	[313]

FORUM OR ISSUER	RECORD	HOW THE PAPER CODES IT, AND ITS POSTURE	DATE	SCORED PRODUCTS NAMED	DOCUMENT	SRC
N.Y. County Supreme Court	Michael Kent v. DoubleVerify Holdings, Inc. et al. (index number masked)	allegation (unverified caption): open per aggregator (14 Oct 2026)	2026-10-08	DoubleVerify (company and directors)	summons and complaint (not read)	[314]
ACCC	ACCC approves Neptune BidCo acquisition of DoubleVerify (Phase 1) (MN-55036)	finding on the merits of a competition review only: assessment completed	2026-10-13	DoubleVerify	regulator determination (competition)	[315]
DoubleVerify (SEC filing)	Fourteen demand letters and two unserved stockholder complaints over the merger (8-K Item 8.01)	party statement (reporting allegations): complaints unserved	2026-10-20	DoubleVerify	issuer filing	[316]
DoubleVerify (SEC filing)	Two merger complaints still unserved (10-Q Q3 2026)	party statement: unserved	2026-10-30	DoubleVerify	issuer filing	[317]

Eight letters and a staff report

Congress asked questions and characterised documents; no agency action on its requests was found ◆ ◆ ◆.

On 7 February 2025, Senators Marsha Blackburn and Richard Blumenthal sent six letters about Adalytics' image-host report to Amazon, Google, DoubleVerify, IAS, MRC and TAG ◆. DoubleVerify's letter asked five questions with answers due on 14 February; the other four recipients' letters were dated from the release without their texts being read [130]. IAS said it answered [79], but no public answer from either verifier was found in the pages opened, without a dedicated search ◆. These letters convey questions and allegations.

Senator Mark Warner's two letters were dated 24 March and released on 28 March 2025 ◆. One asked the FTC to investigate the vendors' bot-filtering claims, with answers due on 28 April; the other sought Justice Department investigation of alleged misrepresentation to federal clients ◆. The FTC letter distinguished general invalid traffic (GIVT) from fraud and used a vendor-linked fraud estimate [132]. The screen found no action on the requests, although Justice Department releases were not separately searched ◆.

Both rounds cite a Marketing Brew report of 11 October 2024, from anonymous sources, that Justice Department and Naval Criminal Investigative Service (NCIS) officials had asked advertising executives about IAS and DoubleVerify; with no primary record found, the item stays unverified ◆.

The House Judiciary Committee's majority staff published "GARM's Harm", an interim report, on 10 July 2024 ◆. It names DoubleVerify, IAS, Oracle Moat, Peer39, Channel Factory, Pixability and Zefr among eleven ad-tech providers GARM invited into its community, and says committee documents suggest GARM was building its framework into their products ◆. It is an investigating committee's staff characterisation, not a finding, and accuses no vendor of unlawful conduct in the passages read; no response was found on Zefr's or Channel Factory's pages ◆ ◆ ◆.

What regulators did

No regulatory act found makes a finding about a scored product's measurement ◆ ◆ ◆ ◆.

The FTC's 20 May 2025 civil investigative demand to the Global Disinformation Index (GDI) named IAS and DoubleVerify among thirteen rating entities ◆. GDI alleged on information and belief that IAS also received a demand; DoubleVerify's proxy repeated news accounts whose description matches IAS only by inference [133]. The FTC withdrew GDI's demand on 16 April 2026 and reported withdrawing other demands ◆. The record does not establish whether either verifier received one.

On 15 April 2026 the FTC and eight states sued Dentsu, WPP Media and Publicis over alleged boycott-related use of the APB and GARM and settled the same day ◆. Dentsu's ten-year order, entered without admission, bars specified political, ideological, journalistic or diversity-related rating criteria unless an advertiser expressly directs their use ◆. No verifier is a defendant. The order governs the agency's reliance on ratings ◆.

In France, Adloox, owned by Peer39 since June 2026, complained in October 2022 that Meta had refused it access to Meta's viewability and brand-safety partner programmes since 2016 ◆ ◆. On 4 May 2023 the Autorité de la concurrence, without deciding the merits, ordered Meta to publish objective, transparent, non-discriminatory and proportionate access criteria; the decision was not appealed, its measures bind Meta and its findings are provisional, and no decision on the merits was listed on 10 November 2026 ◆. The authority sits outside the registers fixed for Q1, as an additional source ◆.

The ACCC's Phase 1 approval of Neptune BidCo US Inc.'s acquisition of DoubleVerify, published on 13 October 2026, is a finding on the merits of a competition review only; the UK Competition and Markets Authority (CMA) and the European Commission returned no case for the vendors queried ◆. The sale, pending on 10 November 2026, belongs to Chapter 7 ◆.

DoubleVerify v. Adalytics

Of the matters open on 10 November 2026, only this suit turns on how a verification product performs; it has one pleading-stage ruling and no judgment ◆. DoubleVerify sued in the District of Maryland on 12 May 2025 over the bot report of 28 March 2025, pleading false advertising under the Lanham Act, defamation, injurious falsehood, tortious interference and unfair competition ◆.

On 27 April 2026 the court dismissed the tortious-interference count without prejudice and let the rest proceed ◆. It held that DoubleVerify had not alleged facts showing that the report's statement that some ads reached bots before the bid was false, but had plausibly alleged that the report falsely implies its pre-bid bot detection does not work, while calling one of DoubleVerify's supports weak ◆. The ruling took DoubleVerify's allegations as true and is not a finding about the report; there is no judgment. Adalytics answered on 11 May 2026, pleading that the statements are true or substantially true and that the case is a SLAPP suit (a strategic lawsuit against public participation) under Maryland law, with no counterclaim ◆. Discovery is under way, and the next status report is due on 9 February 2027 ◆.

The complaint is also DoubleVerify's account of money and accuracy. It says GIVT impressions are deducted from the billable count and sophisticated invalid traffic (SIVT) flagged so advertisers can seek compensation, citing no contract, platform term or

settlement ◆. Its detection figures, such as pre-bid filtering that keeps customers from bidding on or paying for 99% of unwanted GIVT impressions, come with no published test set, denominator or period ◆. Chapter 4 and Chapter 3 test those accounts.

Shareholder suits and filings

Shareholders sued both public verifiers in 2025. The recorded dispositions established no product error rate or factual finding about measurement accuracy ◆ ◆ ◆ ◆.

Electrical Workers Pension Fund, Local 103, I.B.E.W., filed the class action against DoubleVerify in the Southern District of New York on 22 May 2025, for buyers of its stock from 10 November 2023 to 27 February 2025 [134] ◆. Among other things it alleged that DoubleVerify systematically overbilled customers for impressions served to declared bots ◆. The lead plaintiffs dismissed it voluntarily and without prejudice on 22 December 2025, before filing any amended complaint, and no refiled federal action naming DoubleVerify appeared in the party-name search of 10 November 2026 ◆. A voluntary dismissal is neither vindication nor liability. A derivative suit, Kaszirer v. Zagorski, filed against ten officers and directors on 9 December 2025, ended with a notice of voluntary dismissal filed on 4 February 2026 whose prejudice term could not be read ◆.

IAS's class action is about pricing disclosures. Oklahoma Firefighters Pension and Retirement System sued IAS, two officers and Vista Equity Partners on 29 January 2025 over undisclosed price competition from 2 March 2023 to 27 February 2024, after which the stock fell about 41% ◆. On 2 March 2026 the court dismissed it with prejudice for failure to plead a material misrepresentation or omission, without reaching scienter or loss causation; the plaintiff noticed an appeal to the Second Circuit on 26 March, and that it is being briefed rests on the plaintiffs' firm's undated page ◆. Two derivative suits ended in a stipulated voluntary dismissal on 15 January 2026, its terms unread ◆.

IAS's February and May 2025 filings applied a material-adverse-effect test without naming the served class action; it first named the suit in August 2025 ◆. Its later reporting suspension left no periodic filing covering the dismissal or appeal ◆. The 10 November EDGAR screen found no filing containing Adalytics, bounded to that search's coverage ◆. DoubleVerify's FY2025 10-K names the dismissed class action and records $3.5 million in unusual litigation and regulatory costs without allocating them ◆ ◆. Both acknowledge billing disputes, and the federal case-name search found no advertiser or agency suit against either ◆.

By 20 October 2026 DoubleVerify had received fourteen demand letters over its pending sale, and two stockholder complaints alleging deficient proxy disclosures had been filed, none served; it disputes them and names no court or caption ◆. The only caption found, Kent v. DoubleVerify Holdings, filed in New York County Supreme Court on 8 October 2026, is visible only through a commercial aggregator and is not confirmed as one of the two ◆. The buyer is Neptune BidCo US Inc., Nielsen's parent ◆. The author works at Samba TV, which competes with Nielsen in TV measurement.

Other suits

The other suits found concern something other than a verification product, or could not be read ◆. Martin v. Zagorski, filed pro se on 31 March 2025 against DoubleVerify, Providence Equity and four individuals, is a whistleblower-retaliation and civil-rights conspiracy suit, not a product claim; the court dismissed it with prejudice for failure to state a claim on 16 September 2026 ◆. No notice of appeal was on the docket on 10 November, but a motion to extend the time to appeal remains possible until about 18 November ◆.

Chavez v. Human Security, a putative class action filed on 5 May 2026, alleged that HUMAN's website bot-detection code covertly collects visitors' personal data; it names none of HUMAN's ad-fraud products, and the plaintiffs voluntarily dismissed it with prejudice, so ordered on 21 July 2026 ◆. Two older dockets, one against HUMAN reported settled by mediation in May 2024 and one against DoubleVerify dismissed in January 2025, carry no public documents, so their subjects are unknown ◆ [129]. A privacy class action against Comscore, filed in February 2026 and dismissed with leave to amend on 30 June 2026, is not shown to concern its Proximic suitability filters ◆.

GARM in court

The WFA announced GARM's discontinuation three days after X and Rumble sued on 6 August 2024, citing allegations' resource drain ◆ ◆. The announcement gave no court finding of conduct ◆. X's complaint attached a WFA email chain about Meta's Zefr reporting partnership, but the later ruling does not mention Zefr ◆.

The court dismissed X's claims on 26 March 2026, with prejudice for the WFA and domestic defendants for failure to state an antitrust claim, and without prejudice for four foreign defendants on jurisdiction ◆. X appealed; after the WFA's July settlement, their appeal and cross-appeal were dismissed on 11 August [131]. The United States and twelve states filed amicus briefs supporting X's appeal against the others ◆. At the 9 November check, no argument was calendared ◆.

Rumble's parallel suit against the WFA, Diageo, WPP plc and GroupM, filed the same day, was dismissed without prejudice on 13 August 2025 for lack of personal jurisdiction and improper venue; the court did not reach the merits, and no appeal or refiled case had been found by 9 November ◆ ◆. Whether anyone now maintains GARM's floor and tiers is H6, tested in Chapter 8.

Reports and answers

The reports and responses establish different propositions, as Chapter 3 explains ◆ ◆. The Adalytics sitemap screen found six reports naming scored vendors, among fourteen URLs modified since 2023; HUMAN appears in the user-ID report only as a declared partner ◆. Adalytics sells advertiser analytics and states that publishers and ad-tech vendors do not fund its research ◆.

The made-for-advertising (MFA) report, public by 18 March 2024, listed Oracle Moat, IAS, DoubleVerify, HUMAN, Pixalate and Fou Analytics among platforms observed monitoring

ads on MFA sites in January 2024 [◆]. The report of 7 August 2024 showed IAS's publisher-side endpoint returning low-risk values for explicit pages, and ads carrying DoubleVerify or IAS code there, from a convenience sample [◆]. The report of 7 February 2025 relayed advertisers' statements that their vendors had marked all measured impressions on an image host safe or suitable, a second-hand, domain-level account; it showed ads on explicit adult pages there and records, without showing it, one archived instance beside putative abuse material [◆] [◆]. The report of 28 March 2025 showed ads served to bots in crawl datasets, disclaiming magnitude and fault; its statement that advertisers were billed rests on advertisers' non-public log files and invoices [◆] [◆]. Smith and co-authors, at the New Jersey Institute of Technology and DeepSee.io, posted a study on 3 January 2026 measuring how far the publisher-side labels of DoubleVerify, IAS and Oracle agreed, not whether they were right [◆]. Chapter 3 codes what each establishes.

DATE	RESPONDER	ANSWERS	WHAT IT SAYS, PARAPHRASED	EVIDENCE
7 August 2024	DoubleVerify	"AI" brand safety report; also the MFA and Forbes reports	The report lacked clients' campaign settings and its results are manufactured; it had already classified the MFA sites and the Forbes subdomain as MFA	[◆] [◆] [◆]
7 August 2024	IAS	"AI" brand safety report	Almost all tags shown came from its publisher products; the report gave no timeframe or sample	[◆]
10 January 2025	DoubleVerify	Bot report, before publication	Answer in advance on GIVT handling and pre-bid settings	[◆]
7 February 2025	DoubleVerify	Image-host report	Its customers' impressions on the site were 0.000047% of the media transactions it measured in 30 days; pages are classified only above a level of ad traffic; the copy archived that day also put 83% of those impressions on the home page, a figure the live text dropped on 15 October 2025	[◆] [◆]
7 February 2025	IAS	Image-host report	The two domains were 88,000 of nearly 1.9 billion webpages it observed in 60 days; it labelled them ineligible for monetisation	[◆]
13 February 2025	DoubleVerify	Image-host report	Launched a default-on "Highly Illicit: Do Not Monetize" category	[◆]
14 February 2025	IAS	Image-host report and the senators' letter	Says it answered the senators, and that more than 350 of the sites the National Center for Missing & Exploited Children had reported were already covered by its brand-safety products	[79]
21 March 2025	IAS	Bot report, a week before publication	Methodology post on its publisher tools	[◆]
28 March and 3 April 2025	IAS	Bot report	The report tested its publisher tag, not its advertiser tag; it classified URLScan.io impressions as GIVT, SIVT or reduced-value inventory	[◆] [◆]
28 March and 7 April 2025	DoubleVerify	Bot report	GIVT is removed post-bid from billable counts; URLScan.io is not a declared bot; of the report's 115 examples, 62 had no DoubleVerify measurement, 10 fell outside its data retention and all 43 it could match were identified; no per-impression data	[◆]
11 April 2025	MRC	Bot report	Post-serve GIVT filtration is required and pre-bid filtration is not; it has no insight into billing; URLScan.io does not declare itself; it did not test each observation	[◆] [◆]
12 May 2025	DoubleVerify, in its complaint	Bot report	GIVT is deducted from billable counts and SIVT flagged so the advertiser can seek compensation; no contract cited	[◆]
18 December 2025	IAS	No stated trigger; GDI links it to the FTC investigation	Stops using GDI's limited input in Context Control Avoidance, for product reasons	[◆]
11 May 2026	Adalytics, in its answer	DoubleVerify's complaint	Denies liability; pleads truth or substantial truth and Maryland's anti-SLAPP statute	[◆]

Some answers were not found: no public HUMAN response to the bot report in the searches of 10 November 2026, and no vendor response to Smith and co-authors in the vendor pages opened, which was not a dedicated search ◆ ◆. The MRC's answer states its standards and audit scope and says it did not test each observation; the MRC accredits the vendors the report names and had received the senators' February letter ◆ ◆.

Where the search ran

Every absence here is bounded to named registers queried on 10 November 2026, and a zero means only that a register returned nothing for that name ◆ ◆. The sources were fixed at registration: SEC EDGAR; CourtListener and Justia for federal courts; FTC, Justice Department and state attorney-general releases; senate.gov and house.gov; the ACCC, CMA and European Commission registers; the vendors' newsrooms; and Adalytics' and other researchers' publications, searched by vendor and product name. AI research agents ran the searches under the author's direction, and in the verifier pass a separate AI verifier, told to refute each finding, re-ran part of them the same day ◆.

REGISTER	QUERY	RESULT ON 10 NOVEMBER 2026	EVIDENCE
CourtListener, federal party names	DoubleVerify	9 dockets, 5 filed since 1 January 2023	◆
CourtListener, federal party names	Integral Ad Science	7 dockets, 3 since 2023	◆
CourtListener, federal party names	HUMAN Security	3 dockets, 2 since 2023	◆
CourtListener, federal party names	Pixalate	5 dockets, 2 since 2023: the Criteo suit and an unreviewed trademark case that may not name Pixalate	◆ ◆
CourtListener, federal party names	Zefr	None in any year; in a full-text search, the X Corp. exhibits were the only verification matter on the first page	◆ ◆
CourtListener, federal party names	Peer39, Adloox, Scope3, Protected Media (now Protected by Mediaocean), Channel Factory, Fraudlogix, TrafficGuard, Illuma, Pixability	None filed since 2023	◆
CourtListener, federal party names	Adalytics	4 entries since 2023, among them Colossus Media's 2024 suit and DoubleVerify's	[135]
CourtListener, federal party names	Comscore	Not obtained (human-verification page, not bypassed)	◆
SEC EDGAR, full text	"Adalytics", all filers	No filing	◆
UK CMA case finder	DoubleVerify, IAS, HUMAN Security, Pixalate, Zefr, Peer39, Adloox, Protected Media, Channel Factory, Scope3, "brand safety"	No case	◆
European Commission case search	DoubleVerify, Pixalate, Zefr	No case; the IAS query did not render	◆
ACCC acquisitions register	Nielsen and DoubleVerify	MN-55036, approved; determination published 13 October 2026	◆
Autorité de la concurrence decisions	"adloox"; "vérification publicitaire"	Decision 23-MC-01; no later decision on the merits	◆ [136]
Adalytics sitemap	URLs modified in 2023 or later	14 URLs; six reports name scored vendors	◆

CourtListener's access failures prevented completion of the Comscore query and repeat checks of zero results; Justia's Kaszirer docket also remained unread ◆ ◆. The screen

omitted separate Justice Department and state attorney-general searches, party-name queries for mFilterIt, Sightly and FouAnalytics, and publications by FouAnalytics and Check My Ads ◆ ◆. State cases such as Kent and the Delaware IAS-share suit rest on an aggregator or reporting [128] ◆. Sealed and nonpublic matters remain outside the screen. An access failure establishes no absence.

CHAPTER 10

Verdicts in agent buying

The 10 November 2026 screen found verification fields, implemented connectors and interoperability tests in agent buying, while production claims lacked campaign or transaction records tying a verdict to an agent's decision. H7 remains unresolved under the stated reading of its rule. No screened agreement gives such a verdict financial force. The author co-leads the Signals & Measurement working group of AgenticAdvertising.org, which maintains the Ad Context Protocol (AdCP), a declared interest.

IN SHORT

AdCP represents verification thresholds with a named measuring vendor and terms for whose count governs an invoice ◆ ◆. It supplies no measurement accreditation or dispute task in version 3.2 ◆ ◆. IAB Tech Lab's OpenProposal draft represents a viewability guarantee and remedy, without invalid-traffic (IVT) or brand-safety guarantee fields ◆. AAO's 10 November registry held 35 public agents and no measurement-type agent; no agent passed the accepted-measurement-terms test ◆ ◆. DoubleVerify's Butler/Till workflow claim and Zefr's unnamed closed alpha lack a transaction record connecting a verification verdict to an agent's decision ◆ ◆. Trade reporting provides partial corroboration of Butler/Till's quality-analysis work, without settling that link [137, 138]. H7 remains unresolved under the paper's recorded interpretation, and no identified agreement was found giving a verdict financial force ◆.

INTERESTS

The author co-leads AAO's Signals & Measurement working group; findings on AdCP are a declared interest ◆. No Fluff Advisory, the author's company, operates two agents in AAO's registry, and its sales agent is one of the two that pass the billing-window test counted below [6] ◆. The author filed four public comments on OpenProposal on 27 October 2026, none on its measurement fields ◆ [139]. DoubleVerify agreed on 6 August 2026 to be acquired by Neptune BidCo US Inc., the parent of Nielsen, a deal not closed when checked on 10 November 2026, and the author works at Samba TV, which competes with Nielsen in TV measurement ◆ ◆. AdCP's lead architect co-founded Scope3, now Apostra, which founded AAO and lent it seed funding ◆. None of this earns or costs a protocol state.

The test

H7, registered on 9 November 2026, predicts that, by the freeze, verification verdicts appear in agent-buying protocols "at most as schema representation", with no documented production use or demonstrated transaction in an agent-driven buy and none bound to a financial decision by an identified agreement ◆. It covers AdCP, IAB Tech Lab's Agentic Advertising Management Protocols (AAMP) and the scored products' agent interfaces, including servers on the Model Context Protocol (MCP), through which an AI assistant calls outside tools ◆.

Shared rule 12 records five states separately: schema representation, which proves only that a verdict can be expressed; an implemented connector; a demonstrated interoperability test; documented production use; and a demonstrated transaction (Chapter 13). Membership and schema fields earn no deployment credit, and financial force by an identified agreement is a separate field ◆. H7 is supported if "no source screened documents state 4 or state 5 for a verification verdict", contradicted if "a source documents production use (state 4) or a demonstrated transaction (state 5) for a verification verdict", and unresolved if "a production claim rests on a vendor statement with no named buyer, campaign or transaction record (recorded as a vendor assertion)" ◆.

An AI research agent under the author's direction read the sources on 10 November 2026; in the verifier pass, a separate AI agent told to refute each finding re-opened them that afternoon. Neither used a web search engine, so every absence below is bounded to the sources logged ◆. In agentic sources "verification" often means agent or seller identity (AdCP's seller verification, IAB Tech Lab's Agent Registry, HUMAN's Verified AI Agent), which earns no state for an ad-verification verdict ◆.

What AdCP represents

AdCP 3.0.0 introduced vendor-named performance thresholds ◆. Version 3.2.9's subset comprises viewability, IVT, completion rate, brand safety and attention, with IVT a ceiling and the others floors ◆. Measurement terms name whose count governs invoicing, tolerance, billing window and remedies such as credit or extra delivery ◆. These are schema capabilities. The sources screened to 10 November contained no public agreement applying them with a verifier ◆.

AdCP's taxonomy assigns verification to a third party while declining to adjudicate competing vendors ◆. Its limitations page disclaims MRC accreditation and leaves filtration and measurement to the delivery stack or a vendor ◆. The parties can nevertheless choose whose count governs an invoice in the schema ◆.

AdCP 3.2's documentation also stops short of settlement: a structured dispute task "does not exist in 3.2 and is not scheduled for it", a verification vendor's own integration is "out of AdCP scope", and measurement agent interfaces are "a follow-up"; issue #7064, which asks what dispute lifecycle should cover invalid traffic, credits and make-goods, was routed to the Spec Backlog on 21 September 2026 ◆. In 3.2 an ivt or brand_safety standard names the vendor but not which of its metrics is judged: Apostra (formerly Scope3) worked around that with a private extension on 22 October 2026, and the 3.3.0-beta.5 prerelease of 7 November adds an optional metric_id, still a beta on 10 November ◆.

AAO's hosted grader defers measurement-verification assessment until the task surface is frozen ◆. Experimental accreditation fields remain vendor asserted, and the production-evidence program was not selected for 3.3 in the roadmap record read on 10 November ◆ ◆ . No protocol-body production matrix for verification verdicts was found there ◆.

What AAMP represents

IAB Tech Lab develops all its agentic work under AAMP, and AAMP represents less than AdCP ◆. OpenProposal, published as AAMP 3.0 draft 1 for comment on 22 September 2026 and not a ratified standard, represents a committed viewable_rate on an MRC basis, a selectable measurement_source vendor (its example offers IAS and DoubleVerify) and a shortfall_remedy of make-good or credit, with no invalid-traffic or brand-safety guarantee field ◆. Tech Lab's release page extends comment to 10 November 2026, while the repository still shows 22 October ◆ [140]. None of the three fields was on the main branches of the AAMP buyer-agent, seller-agent or primitives repositories on 10 November; an unmerged seller-agent branch of 19 October models them and maps Google Ad Manager's viewability partner to measurement_source ◆ [141].

The Agentic Real Time Framework (ARTF) 1.0 of 12 November 2025 names pre-impression fraud detection as a use case and lets an agent attach metric objects from OpenRTB, the real-time bidding specification, to an impression before the auction ◆. Those are predictions used to select inventory, not post-delivery verdicts, and the reference implementation shows them in a sample with no verification vendor's connector ◆. The AAMP buyer agent (2.4.2) models brand safety as excluded content categories and keywords, and its example brief lists IAS and DoubleVerify as required vendors in a field no code reads ◆.

What the vendors have built

The scored vendors' agent interfaces are mostly their own tools, and where one is said to reach a buying protocol, the vendor is the only source ◆.

DATE	VENDOR	INTERFACE OR CLAIM	STATE REACHED, AS CODED	EVIDENCE
21 Jul 2025	HUMAN	Verified AI Agent: open-source demonstration of agent identity	None for an IVT verdict (agent identity)	◆
24 Jul 2025	HUMAN	Open-source MCP server for its security products	None for an IVT verdict	◆
4 Nov 2025	DoubleVerify	DV AI Verification, aimed at AI-agent traffic and AI-generated content	No protocol state	◆
16 Dec 2025	IAS	IAS Agent, an assistant inside IAS's interface; dashboard feature launched March 2026	No MCP, AdCP or AAMP interface found	◆
30 Apr 2026	Zefr	Zain: AdCP-based YouTube buying integration and an MCP server	Connector, as the vendor states it	◆
16 Jun 2026	IAS	Launch partner of Databricks CustomerLake, signals to agents through clean rooms	No protocol state	[142]
17 Jun 2026	DoubleVerify	DV Neura: Insight Agent over MCP "today"; Activation Agent forecast for the third quarter	Connector, as the vendor states it	◆
24 Jul 2026	DoubleVerify	Blog: Butler/Till buys through AdCP; PubMatic applies DoubleVerify protections before the auction	Production use, vendor assertion, uncorroborated	◆

DATE	VENDOR	INTERFACE OR CLAIM	STATE REACHED, AS CODED	EVIDENCE
14 Sep 2026	Zefr	Closed alpha with unnamed partners; Rapid Response Blocking Agent	Production use, vendor assertion, no partner named	◆
17 Sep 2026	DoubleVerify	Activation Agent "expected to become available in the months ahead"	Forecast	◆
5 Oct 2026	IAS; DoubleVerify	ChatGPT Ads pilots (measurement; suitability evaluation)	No protocol state	◆
Undated, read 10 Nov 2026	Pixalate	Analytics and Enrichment MCP connectors and IQ Agent for a client's own IVT data; OpenEPG MCP planned	Connector for retrieving data, as documented by the vendor	◆
Read 10 Nov 2026	Channel Factory	AERO agents for lists, bidding and reports; no protocol named	No protocol state	◆
Read 10 Nov 2026	Peer39 and Adloox	No agent interface found; statements of intent	None found	◆ ◆

DoubleVerify's June release claimed an available Insight Agent over MCP and forecast an Activation Agent; its September release still forecast the latter ◆. IAS Agent runs inside IAS's interface, and HUMAN's public MCP and identity tools establish no ad-verdict integration in the sources screened ◆ ◆. Those records are kept apart from production use.

Zefr's Zain pages claim an AdCP buying integration, an MCP server and a closed alpha, without documenting a suitability verdict in an AdCP field ◆ ◆. Pixalate documents connectors for querying client IVT data without an AdCP or AAMP implementation ◆. Channel Factory names no protocol for AERO, while Peer39 and Adloox supplied statements of intent in the screened sources ◆ ◆ ◆.

For Protected by Mediaocean (formerly Protected Media), Proximic by Comscore, Illuma, Fraudlogix and Escalated the screen went no further than AAO's member list, where none has a public profile, and, for Protected by Mediaocean, a brand.json check that found none ◆ ◆. For mFilterIt, FouAnalytics, Pixability and Sightly the stream kept no record, so this chapter bounds no absence for them. IAS's closed pilot and DoubleVerify's pilot in development, both announced for ChatGPT Ads on 5 October 2026, are third-party measurement in a channel that AdCP's taxonomy models as a closed loop in which the seller is also the measurement vendor, first-party validation in this paper's terms (Chapter 1); neither is an agent-buying protocol integration ◆.

The instrument below dates every release, announcement, pilot and claim, keeping announcement, availability and forecast apart.

Verification in agent buying, dated

Protocol releases and vendor announcements that concern verification verdicts in agent-driven buying. Status is as stated by the publisher; "Independently observable" means a public specification, document or code, not use in a transaction.

Machine-readable copy: data/agentic-map.csv.

PUBLISHER OR VENDOR	EVENT	STATUS	DATE	INDEPENDENTLY OBSERVABLE?	WHAT THE RECORD SUPPORTS	SRC
HUMAN	Verified AI Agent open-sourced (RFC 9421 agent identity)	LIVE	2025-07-21	Yes: public code	Published or released; for a vendor feature, availability as the vendor states it.	[341]
HUMAN	Open-source MCP server for Sightline Cyberfraud Defense and Client-side Defense	LIVE	2025-07-24	Yes: public code	Published or released; for a vendor feature, availability as the vendor states it.	[340]
DoubleVerify	DV AI Verification announced (AI-agent traffic); GA said to be 'later this year'	ANNOUNCED	2025-11-04	No: the vendor's own statement	Announced or proposed; not shown in use.	[348]
IAB Tech Lab	Agentic Real Time Framework v1.0 released	LIVE	2025-11-12	Yes: a public specification or document	Published or released; for a vendor feature, availability as the vendor states it.	[331]
PubMatic / Butler/Till	AdCP-enabled CTV campaign for Clubtails announced; no verification vendor named	UNCONFIRMED	2025-12-15	No: the vendor's own statement	A vendor's claim of use with no named campaign record or transaction (coded a vendor assertion for H7).	[349]
IAS	IAS Agent announced; global rollout expected early Q1 2026	ANNOUNCED	2025-12-16	No: the vendor's own statement	Announced or proposed; not shown in use.	[336]
IAS	IAS Agent Overview Dashboard Insights launched (as stated 2026-06-05)	LIVE	2026-03	No: the vendor's own statement	Published or released; for a vendor feature, availability as the vendor states it.	[337]
AgenticAdvertising.org	AdCP 3.0.0 released with performance standards (viewability, ivt, brand_safety), measurement terms, makegood remedies and content standards	LIVE	2026-04-22	Yes: a public specification or document	Published or released; for a vendor feature, availability as the vendor states it.	[350] [320]
Zefr	Zain agentic hub, AdCP-based YouTube integration and MCP server announced	ANNOUNCED	2026-05-01	No: the vendor's own statement	Announced or proposed; not shown in use.	[343]
AgenticAdvertising.org	Measurement Taxonomy page added (PR #3843)	LIVE	2026-05-02	Yes: a public specification or document	Published or released; for a vendor feature, availability as the vendor states it.	[350] [318]
DoubleVerify	DV Neura announced; Insight Agent via MCP 'today'; Activation Agent forecast for Q3	ANNOUNCED	2026-06-17	No: the vendor's own statement	Announced or proposed; not shown in use.	[333]
AgenticAdvertising.org	AdCP 3.1.0 released; experimental measurement capability block	LIVE	2026-06-18	Yes: a public specification or document	Published or released; for a vendor feature, availability as the vendor states it.	[350] [320]
DoubleVerify	Blog says Butler/Till buys via AdCP with DV protections applied pre-auction by PubMatic AgenticOS	UNCONFIRMED	2026-07-24	No: the vendor's own statement	A vendor's claim of use with no named campaign record or transaction (coded a vendor assertion for H7).	[323]
IAB Tech Lab	AAMP 2.3: vendor approval gate (off by default), pricing provenance	LIVE	2026-07-30	Yes: a public specification or document	Published or released; for a vendor feature, availability as the vendor states it.	[332]

PUBLISHER OR VENDOR	EVENT	STATUS	DATE	INDEPENDENTLY OBSERVABLE?	WHAT THE RECORD SUPPORTS	SRC
PubMatic	Governance architecture for AgenticOS; 80+ autonomous campaigns, 4,000+ direct deals claimed	ANNOUNCED	2026-08-05	No: the vendor's own statement	Announced or proposed; not shown in use.	[351]
AgenticAdvertising.org	Issues #7064 (commercial lifecycle incl. IVT disputes, makegoods) and #7068 (production evidence) opened	ANNOUNCED	2026-08-30	Yes: a public specification or document	Announced or proposed; not shown in use.	[352] [353]
Zefr	Closed alpha of agentic YouTube activation; Rapid Response Blocking Agent	BETA	2026-09-14	No: the vendor's own statement	Limited, pilot or pre-release stage as stated.	[344]
DoubleVerify	Ask Neura launched; Activation Agent 'expected ... in the months ahead'	LIVE	2026-09-17	No: the vendor's own statement	Published or released; for a vendor feature, availability as the vendor states it.	[335]
IAB Tech Lab	AAMP 3.0 / OpenProposal draft 1 published for comment (to 2026-10-22)	BETA	2026-09-22	Yes: a public specification or document	Limited, pilot or pre-release stage as stated.	[328] [327]
AgenticAdvertising.org	AdCP 3.2.1 GA release object published (docs say 2026-09-30)	LIVE	2026-10-01	Yes: a public specification or document	Published or released; for a vendor feature, availability as the vendor states it.	[350] [354]
IAS; DoubleVerify	ChatGPT Ads brand safety measurement in closed pilot (IAS); suitability evaluation pilot in development with OpenAI (DV)	BETA	2026-10-05	No: the vendor's own statement	Limited, pilot or pre-release stage as stated.	[355] [356]
Apostra	Interchange scores ivt and brand_safety standards via a private metric_id extension	LIVE	2026-10-22	Yes: a public specification or document	Published or released; for a vendor feature, availability as the vendor states it.	[321]
AgenticAdvertising.org	AdCP 3.3.0-beta.5 adds optional metric_id to performance standards	BETA	2026-11-07	Yes: a public specification or document	Limited, pilot or pre-release stage as stated.	[357] [350]
AgenticAdvertising.org	AdCP 3.2.9 released (latest GA-line patch)	LIVE	2026-11-09	Yes: a public specification or document	Published or released; for a vendor feature, availability as the vendor states it.	[350]

The registry, as what it is

AAO's registry lists the agents its members choose to register, and its conformance tests run on AAO's hosted runner against fictional test kits: it records what registered agents passed in a test, not what anyone bought ◆ ◆. On 10 November 2026 the public listing held 35 agents and none of type measurement or governance; AAO's member records add 24 members-only and 19 private agent entries, none a verification vendor's, and an unregistered agent would not appear at all ◆.

Vastlint and TVIQ's agent each passed ten content-standards test steps, establishing implementations and tests for creative verdicts rather than a scored verifier ◆. No agent passed accepted measurement terms, and one of the two billing-window passes belongs to the author's firm ◆. The table retains every test denominator.

AAO'S PUBLIC RECORDS, READ 10 NOVEMBER 2026	VALUE	EVIDENCE
Agents listed (sales / signals / creative / buying / unknown)	35 (28 / 4 / 1 / 1 / 1)	◆
Agents of type measurement or governance in the public listing	0	◆
Agents exposing content-standards tools; passing content_standards	2; 2 (10 of 10 steps)	◆
measurement_terms_accepted: passing / untested / no record	0 / 11 / 24	◆
measurement_terms_rejected: passing	9	◆
billing_finality_delivery: passing / partial / untested / no record (one pass is the author's company's agent)	2 / 5 / 11 / 17	◆
vendor_metric_accountability: passing / untested / no record (fictional attention vendor)	3 / 15 / 17	◆
Public member profiles	159	◆
Verification-vendor domains serving a brand.json with agents (12 checked)	0	◆
AdCP compliance specialisms at 3.3.0-beta.5; for measurement or verification	32; 0	◆

The two production claims

Two sources claim production use of a verification verdict in agent buying, and both are the vendors' own ◆ ◆. DoubleVerify's blog of 24 July 2026 says that "Today, Butler/Till is using Claude as its buyer agent" to activate media buys through AdCP, connected to DoubleVerify's platform over MCP, while PubMatic's AgenticOS applies the selected DoubleVerify protections to inventory packages before the auction ◆. It names the agency but no advertiser, campaign, date range, spend or transaction record, and the agency's own words, found only in DoubleVerify's June release, name no buy or campaign ◆. Taken at its word, it describes pre-bid exclusion, which changes what is bought, not what is paid (Chapter 4), and it does not say DoubleVerify's criteria travel in an AdCP field ◆.

The original source screen did not corroborate the specific DoubleVerify–AdCP workflow ◆. An expanded search on 11 November found Digiday's 24 June report that Butler/Till was developing Claude agents with DoubleVerify to analyse inventory and steer campaigns away from poor-quality placements [137]. Digiday's March report names Geloso's December 2025–January 2026 flights and compares quality with a DoubleVerify benchmark; Marketing Dive and MediaPost relay the parties' claim that less than 1% failed DoubleVerify standards [138, 143, 144]. These reports add campaign and buyer context, but do not establish

that an actual DoubleVerify verdict determined the agent's purchase, which metric was used, or whether it travelled in the protocol. The later DV workflow could describe a different buy. The sources are counterevidence to a blanket absence of other-party reporting, not a transaction record settling H7.

Zefr's alpha is the second claim: its page says partners are "activating YouTube campaigns through Zain", which suggests live campaigns, but it names no partner, campaign or transaction and does not say Zefr's suitability exclusions apply in them ◆.

H7, coded by the rule

H7 remains unresolved under the paper's recorded interpretation ◆. The source record establishes schemas, connectors and tests; the phrase in H7's statement that use stops at schema representation is already too broad ◆ ◆. Production requires the missing connection between a verdict and a live decision. Newly found buyer reporting narrows the evidentiary gap but does not by itself supply that connection [137, 138]. No identified financial-authority agreement was found in the screened record ◆.

The figure sets out the coded states for each protocol, verdict type and vendor interface, with the agreement field beside them ◆.

VERIFICATION VERDICTS IN AGENT BUYING: THE FIVE PROTOCOL STATES AND FINANCIAL FORCE

	Schema	Connector	Interop test	Production use	Transaction	Financial force
AdCP: viewability	yes	yes				draft/schema
AdCP: invalid traffic	yes	yes				draft/schema
AdCP: brand safety rate	yes	yes				draft/schema
AdCP content standards (non-verifier tes	yes	yes	yes			
AdCP: attention score	yes					draft/schema
AAMP OpenProposal: viewability guarantee	draft/schema					draft/schema
AAMP ARTF: pre-impression fraud and view	yes	yes				
AAMP buyer and seller agents: brand safe	draft/schema					
DoubleVerify DV Neura: pre-bid brand sui	n/a	vendor says		vendor says		
IAS Agent: brand safety and suitability						
HUMAN: ad fraud						
Zefr Zain: brand suitability exclusions	n/a	vendor says		vendor says		
Apostra Interchange: IVT and brand safet	yes	yes				
Peer39/Adloox, Channel Factory: all						
Pixalate: Analytics and Enrichment		vendor says				

Documented · Vendor assertion only · Draft or schema field only · None found · Not applicable

The author co-leads AdCP's Signals & Measurement working group, a declared interest. AdCP's own measurement taxonomy says the protocol does not adjudicate between verification vendors. Read 10 November 2026.

> **Verdicts are schema fields in agent buying; no transaction and no binding agreement were found.** For each protocol and vendor agent interface, the five protocol states of shared rule 12 (schema representation, implemented connector, interoperability test, production use, demonstrated transaction) and whether an identified agreement gives a verification verdict financial force. The only production claims are vendor assertions.
>
> Units: coded state · Scope: AdCP, AAMP and the scored vendors' agent interfaces, to 10 November 2026 · Sources [318] [319] [320] [6] [321] [322] [323] [324] [325] [326] [327] [328] [329] [330] [331] [332] [333] [334] [335] [336] [337] [338] [339] [340] [341] [342] [343] [344] [345] [346] [347]
>
> Membership and schema fields earn no deployment credit; the states are status-sensitive.

The rule is ambiguous because DoubleVerify names an agency buyer, while its unresolved clause refers to a vendor claim without a named buyer, campaign or transaction record ◆. The primary reading requires more than that agency name to establish a verdict in production. The author reads it that way: an agency name without a verdict-to-decision record is still a vendor assertion. Zefr still names no alpha partner ◆.

What would settle it

H7 would be contradicted by a documented live buy showing which verification verdict changed the agent's action, with a named campaign or transaction and its dates ◆ ◆. A delivered and billed AdCP package naming the governing verifier would be one route. Financial authority additionally requires the identified agreement making that count determine an invoice or remedy ◆.

The same parties could settle it the other way. If Butler/Till, PubMatic or DoubleVerify described the agent-driven buys without a third-party verdict applied, or Zefr's alpha proved to carry no suitability verdict, and no other source documented state 4 or 5, H7 would be supported as registered ◆. The recheck at the freeze, on 11 December 2026, covers AAO's registry counts, the brand.json files, AdCP 3.3's changes, the disposition of OpenProposal's comments and the availability of DoubleVerify's Activation Agent ◆ ◆ ◆ ◆.

CHAPTER 11

The Verdict Quotient

The Verdict Quotient (VQ) summarises the public record of a product's documented capability within its buying task. Independent detection evidence is the lowest-scoring dimension in the suite and invalid-traffic cohorts; accreditation is lowest among suitability specialists. Those are gaps in the evidence a buyer can inspect. They do not measure which service makes the fewest errors.

> **IN SHORT**
>
> The matrix scores 17 products on 12 dimensions, subject to each cohort's applicability mask. The field VQ is 60 for suites, 53 for invalid-traffic specialists and 39 for suitability specialists; these are separate reference points, not a comparison between cohorts. Only cohort C meets the minimum size for bands. It has 0 products in Most documented, 3 in Well documented, 3 in Partly documented and 1 in Least documented. Of its banded rows, 4 retain their band under every tested reweighting. Missing-evidence scenarios change 6 bands. The quotient describes disclosure and documented capability, without establishing detection accuracy.

What the quotient does

The dimensions give each attribute one home. Depth covers method disclosure, independent detection evidence, invalid-traffic method, safety and suitability method, and verdict uses and authority. Reach covers environment coverage, platform access, measured scale and US availability and use. Standing covers accreditation, interest disclosure and safeguards, and the dispute and correction process. Each dimension has published anchors from 0 to 5; the applicable set depends on the product's cohort. Chapter 13 publishes the rules and their limits.

The group means use counted cells only. An n/e cell means evidence was not established; n/a means the dimension falls outside the cohort's mask. Neither becomes zero. Raw = 0.45 Depth + 0.25 Reach + 0.30 Standing, with weights renormalised if a group is absent. The separate coverage field shows how much evidence underlies that average. A fuller record can lower an average when the additional evidence supports a low anchor.

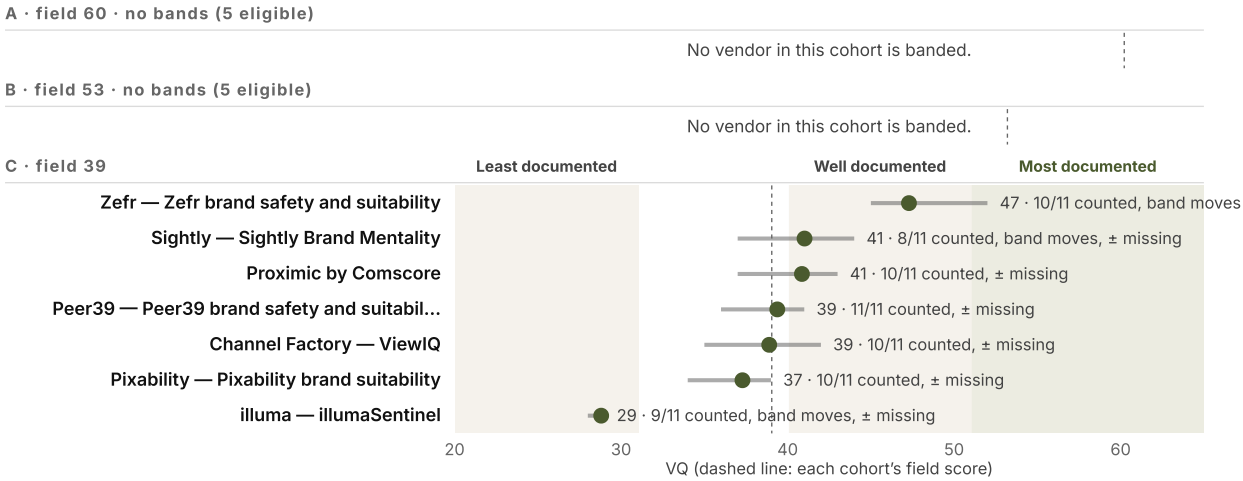
The calculation pulls Raw toward the cohort field mean: Shrunk = (k × Raw + K × F) / (k + K). Here k is the row's counted-cell total, F is the mean Raw of admitted cohort rows, and K is 4. VQ = 100 × Shrunk / 5. K is the author's heuristic for moderating thin records, not a parameter estimated from observed error rates. Shrinkage leaves missing cells missing; it supplies no evidence about them.

Bands use the unrounded offset 20 × (Shrunk − F). Most documented begins at +12, Well documented at +1, Partly documented at −8, and Least documented below −8. Equal offsets

receive the same band. Display rounding does not decide membership. A row needs 6 counted cells with every applicable group represented; a cohort needs 6 eligible rows before any bands appear. Cohorts A and B each contain 5 rows and are therefore shown without bands. The cohort gate applies even when an individual row has a complete record.

The results by buying task

VERDICT QUOTIENT (0–100) BY COHORT, WITH THE RANGE UNDER FOUR REWEIGHTINGS



Grey bar: range under the four reweightings. Not plotted, A to Z: Adloox — Adloox verification (cohort too small), DoubleVerify — DoubleVerify advertiser ve... (cohort too small), Integral Ad Science — IAS advertiser veri... (cohort too small), mFilterIt — mFilterIt Ad Traffic Validati... (cohort too small), Protected by Mediaocean (cohort too small), Escalated — Escalated IVT scoring (cohort too small), FouAnalytics (cohort too small), Fraudlogix — Fraudlogix IVT detection (cohort too small), HUMAN — HUMAN Ad Integrity Suite (cohort too small), Pixalate — Pixalate ad fraud protection a... (cohort too small). Their scores are in data/vq.csv.

> **Read the band, not the rank: no product clears the most documented line, and ten rows are not banded.** The Verdict Quotient for the products with enough evidence to band, one section per cohort. The Verdict Quotient uses the band offsets of the author's Contextual Quotient, but its field mean is the cohort's mean Raw score rather than a pooled group mean, so the two do not read identically; neither construction validates it. Depth, Reach, Standing group means are combined 45/25/30 and shrunk toward the cohort's field mean (the mean Raw score of its admitted rows) with a prior strength of 4; bands come from the unrounded distance to that field score, and only verified or qualified cells count. A band needs at least 6 counted cells with one in every applicable group, in a cohort with at least 6 such rows. Across the reweightings, positions move by up to 1 places within a cohort, and 4 of the 7 banded products keep their band under every weighting. VQ measures publicly documented capability within a cohort, not product performance.
>
> Units: VQ points on a 0–100 scale, pulled toward the cohort field mean; shown rounded, plotted unrounded · Scope: 7 banded products of 17 in the capability matrix, scores from evidence read in November 2026; recheck on 11 December 2026
>
> Single scorer. The index inherits every limit of the dimension scores and is sensitive to the weights; the bands, not the positions, are the finding.

Cohort A contains full verification suites, cohort B impression-level invalid-traffic (IVT) specialists, and cohort C separately purchased suitability services. The field values are 60, 53 and 39, respectively. A field value is a reference computed from that cohort's admitted rows. It is neither a market average weighted by spend nor a benchmark for a different

buying task. A suite's broader mask must not turn into a claim that it outperforms a specialist.

In cohort C, the displayed lines are 51 for Most documented, 40 for Well documented and 31.1 for Partly documented. The calculation applies the full unrounded lines. No row reaches Most documented under the primary settings. The table groups products by band and uses alphabetical order within each band, avoiding a league table.

COHORT C PRODUCT	PRIMARY BAND	SAME BAND UNDER ALL REWEIGHTINGS	MISSING-EVIDENCE FLAG
Proximic by Comscore	Well documented	Yes	±
Sightly Brand Mentality	Well documented	No	±
Zefr brand safety and suitability	Well documented	No	None
Channel Factory ViewIQ	Partly documented	Yes	±
Peer39 brand safety and suitability	Partly documented	Yes	±
Pixability brand suitability	Partly documented	Yes	±
illumaSentinel	Least documented	No	±

The flags answer different questions. Reweighting stability asks whether changing the importance of the groups moves a band. A published sensitivity flag must cover a band change or eligibility loss under any single-cell deletion, as well as a band change under either missing-evidence scenario. The table separates these tests; six of the seven cohort C rows are deletion-sensitive under the current scores (removing Peer39's detection-evidence cell moves it from 39.38 to 42.05, and Pixability's accreditation cell from 37.29 to 40.11, both into Well documented), four are scenario-sensitive, and Zefr is affected by neither. Neither test is a confidence interval. A row can keep its band under every weighting and still carry ±, as Proximic, Channel Factory, Peer39 and Pixability do; Sightly and illuma also move under at least one weighting. The matrix exposes the underlying cells and their reasons rather than asking a reader to rely on the label alone.

The capability matrix

Each cell scores one documented capability from 0 to 5 against fixed anchors; n/e means no evidence, n/a means the dimension does not apply to the cohort. Click a cell for its rationale, availability, evidence grade, date, sources and verification status. Rows are compared only within their cohort, and a score is documented capability, not accuracy.

5 documented in operational detail · 3 real but thin or narrow · 1 adjacent or incidental

0 documented absence · n/e insufficient evidence · n/a outside the cohort's mask

unverified: shown, counted as n/e

β beta · ann announced · ? unconfirmed · ret retired · ± provisional band: it moves when missing evidence is filled

Product	VQ	Method	Detection	IVT	Suitability	Authority	Environments	Platforms	Scale	US	Accredited	Interests	Disputes	Counted
A · FULL VERIFICATION SUITES														
NO BANDS IN THIS COHORT, A TO Z														
Adloox ADLOOX VERIFICATION · VERIFICATION · OPERATING; ACQUIRED BY PEER39 (JUNE 2026, COMPLETION RECORDED IN A UK FILING); BRAND AND SERVICE RETAINED; MRC ACCREDITATION CONTINUED ON 29 JUNE 2026	57 COHORT TOO SMALL	3	1	3	3	4	4	2	2	4	3	3	2	12/12
DoubleVerify ◆ DOUBLEVERIFY ADVERTISER VERIFICATION · VERIFICATION · ACTIVE, LISTED; MERGER WITH A SUBSIDIARY OF NEPTUNE BIDCO US INC. (NIELSEN'S PARENT) PENDING AT 10 NOVEMBER 2026	68 COHORT TOO SMALL	3	2	3	3	4	4	5	5	5	3	4	3	12/12
Integral Ad Science IAS ADVERTISER VERIFICATION · VERIFICATION · ACTIVE; PRIVATELY HELD PORTFOLIO COMPANY OF NOVACAP-MANAGED FUNDS; NO LONGER FILES PERIODIC REPORTS	66 COHORT TOO SMALL	3	1	3	3	4	4	5	n/e	5	3	4	3	11/12
mFilterIt MFILTERIT AD TRAFFIC VALIDATION · VERIFICATION	49 COHORT TOO SMALL	3	1	3	3	2	3	n/e	3	3	0	3	1	11/12
Protected by Mediaocean VERIFICATION	61 COHORT TOO SMALL	3	1	3	3	4	4	n/e	n/e	5	4	3	1	9/12
B · INVALID-TRAFFIC AND FRAUD SPECIALISTS														
NO BANDS IN THIS COHORT, A TO Z														
Escalated ESCALATED IVT SCORING · FRAUD	49 COHORT TOO SMALL	3	2	3	n/a	2	3	n/a	n/e	3	1	3	1	9/10
FouAnalytics FRAUD	46 COHORT TOO SMALL	3	1	3	n/a	1	3	n/a	n/e	4	0	3	1	9/10

Product	VQ	Method	Detection	IVT	Suitability	Authority	Environments	Platforms	Scale	US	Accredited	Interests	Disputes	Counted
Fraudlogix FRAUDLOGIX IVT DETECTION · FRAUD	46 COHORT TOO SMALL	1	1	3	n/a	2	3	n/a	3	3	0	3	3	10/10
HUMAN HUMAN AD INTEGRITY SUITE · FRAUD · ACTIVE; PRIVATE	63 COHORT TOO SMALL	3	1	3	n/a	4	4	n/a	n/e	5	4	3	3	9/10
Pixalate PIXALATE AD FRAUD PROTECTION AND ANALYTICS · FRAUD · ACTIVE; PRIVATE	62 COHORT TOO SMALL	3	1	4	n/a	4	4	n/a	3	5	3	3	3	10/10
C · CONTENT SAFETY AND SUITABILITY SPECIALISTS														
BANDED, BY VQ														
Zefr ZEFR BRAND SAFETY AND SUITABILITY · CONTEXTUAL · ACTIVE, PRIVATE	**47** WELL DOCUMENTE	3	1	n/a	3	2	2	4	n/e	5	2	3	1	10/11
Sightly SIGHTLY BRAND MENTALITY · CONTEXTUAL	**41**± WELL DOCUMENTE	3	n/e	n/a	3	2	1	n/e	n/e	3	0	3	1	8/11
Proximic by Comscore CONTEXTUAL	**41**± WELL DOCUMENTE	3	1	n/a	3	2	4	1	n/e	3	0	3	1	10/11
Peer39 PEER39 BRAND SAFETY AND SUITABILITY · CONTEXTUAL	**39**± PARTLY DOCUMENTE	3	1	n/a	3	2	3	1	2	3	0	3	1	11/11
Channel Factory VIEWIQ · CONTEXTUAL · ACTIVE; PRIVATE	**39**± PARTLY DOCUMENTE	3	2	n/a	3	2	1 ?	1	n/e	3	0	3	1	10/11
Pixability PIXABILITY BRAND SUITABILITY · CONTEXTUAL	**37**± PARTLY DOCUMENTE	3	1	n/a	3	2	1 ?	1	n/e	3	0	3	1	10/11
illuma ILLUMASENTINEL · CONTEXTUAL	**29**± LEAST DOCUMENTE	1	1	n/a	1	2	1	n/e	n/e	1	0	3	1	9/11

How VQ is built. The Verdict Quotient uses the band offsets of the author's Contextual Quotient, but its field mean is the cohort's mean Raw score rather than a pooled group mean, so the two do not read identically; neither construction validates it. Three groups: **Depth** (method disclosure, independent detection evidence, invalid-traffic method, safety and suitability method, verdict uses and authority), **Reach** (environment coverage, platform access, measured scale, us availability and use) and **Standing** (accreditation, interest disclosure and safeguards, dispute and correction process). Only verified or qualified cells count; an unverified cell counts as n/e, and n/a comes only from the cohort's applicability mask. Each group is the mean of the counted cells inside it, so an unscored dimension is not counted as a zero. Raw combines the groups at 45/25/30, renormalised over the groups present, and is shrunk toward the cohort's field mean, the mean Raw score of the cohort's admitted rows, thin rows included, with a prior strength of 4 against the

row's own count of counted cells. A cohort is a set of products doing a comparable job; each has its own field mean and band lines, and rows are never ranked across cohorts. The weights and the prior strength are author-defined heuristics. Bands come from the unrounded distance to the cohort's field score (Most documented at +12 or more, Well documented at +1 or more, Partly documented at −8 or more, Least documented below that); only the displayed score is rounded. A row gets a band only with at least six counted cells and at least one in every applicable group, and only in a cohort with at least 6 such rows; the others show a score without a band and are listed apart, A to Z, with any missing group named. The three bars under each score are Depth, Reach and Standing in that order. VQ summarises publicly documented capability within a cohort, not product performance, and is not calibrated to buyer outcomes.

Read the band, not the rank. A · Full verification suites (5 rows, field VQ 60.2). Only 5 rows are band-eligible, under the 6 a cohort needs, so its scores and coverage are shown without bands. Deleting one counted cell at a time moves a score by up to 3 points up (Protected by Mediaocean, dispute and correction process) or 4 down (mFilterIt, interest disclosure and safeguards); 0 deletions would drop a row below admission and 0 would cost a row its band eligibility. **B · Invalid-traffic and fraud specialists** (5 rows, field VQ 53.2). Only 5 rows are band-eligible, under the 6 a cohort needs, so its scores and coverage are shown without bands. Deleting one counted cell at a time moves a score by up to 4 points up (Fraudlogix, accreditation) or 3 down (FouAnalytics, interest disclosure and safeguards); 0 deletions would drop a row below admission and 0 would cost a row its band eligibility. **C · Content safety and suitability specialists** (7 rows, field VQ 39.0). Band lines: Most documented 51.1, Well documented 40.1, Partly documented 31.1, Least documented below. Under four reweightings (equal thirds, depth-led, reach-led, standing-led), with Raw and the field mean recomputed, no product sits in the top band (Most documented) (Zefr joins it under reach-led weights), while positions move by up to 1 places and by 0.3 on average. Of the 7 banded products, 4 keep their band under every weighting (Channel Factory, Peer39, Pixability, Proximic by Comscore); illuma, Sightly, Zefr cross a band line under at least one. With a prior strength of 0 (Raw, no shrinkage), 2 and 8 instead of 4, positions move by up to 0, 0 and 0 places. Deleting one counted cell at a time moves a score by up to 3 points up (Sightly, environment coverage) or 3 down (Peer39, interest disclosure and safeguards); 0 deletions would drop a row below admission and 0 would cost a row its band eligibility. In the missing-evidence scenarios (every n/e cell of a banded row filled with 1, then with 4, with the field and every other row frozen), 6 of the 6 banded rows with n/e cells change band (Channel Factory, illuma, Peer39, Pixability, Proximic by Comscore, Sightly; marked ±). **Missing evidence is not neutral:** every deletion case is in data/vq-missing-evidence.json, and every row's scenario scores, including a hypothetical 0-to-5 envelope reported separately and never used for the flag, are in data/vq.csv. These tests cover weights, prior strength and missing cells, not evidence grades, the scorer or the product sample. Every input is in the CSVs in the appendix, so the index can be recomputed on different weights.

Mantis Sentinel was listed rather than scored after the final review: the scored-set rule requires a product sold apart from media and live or in beta, and its advertiser page establishes neither. Dropping it raised the cohort C field mean and moved Channel Factory and Peer39 from Well documented to Partly documented.

Where the field is weak and strong

Computed across numeric cells in the supplied matrix, independent detection evidence has the lowest dimension mean in A and B. US availability and use has the highest in each. Within A, platform access and environment coverage are also relatively well documented; the dispute process is weak. Within B, accreditation follows detection evidence among the weakest dimensions, while environment coverage follows US use among the strongest. A

missing cell is excluded from each dimension mean, so the dimensions can have different contributing populations.

In cohort C, accreditation has the lowest mean, followed by dispute and correction process, then independent detection evidence. Interest disclosure and safeguards has the highest mean, followed by US availability. The interest cells credit disclosure of sell-side relationships. They do not establish that those relationships are harmless or that the vendor has a complete separation policy. Each score means only what its anchor says, which prevents a familiar label such as trust from swallowing distinct questions.

The low detection cells fit the bounded study result: no independent test of a scored product's sold category with a labelled test set and published denominator was found in the sources screened to 10 November 2026 ◆. The low accreditation cells have a different basis, the match between documented sold units and MRC's scoped record ◆. An error-rate test and an accreditation letter answer separate questions. Neither the quotient nor its strongest dimension supplies the missing answer.

How stable the bands are

The build recomputes the index under equal group weights, Depth-led 60/20/20, Reach-led 30/45/25 and Standing-led 30/20/50, recalculating Raw and F. In cohort C, 4 of 7 rows keep their band under all variants. Zefr reaches Most documented only under the Reach-led alternative. Sightly falls to Partly documented under the Reach-led and Standing-led alternatives, and illuma rises to Partly documented under the Standing-led one. These are consequences of the chosen weights, not new observations of the products.

The supplied reweighting statistic reports maximum within-cohort position changes of 0 in A, 1 in B and 1 in C. That diagnostic tests numerical movement; the paper makes no product ranking from it. In A and B, zero band changes merely reflect the absence of bands. They cannot establish stability of a distinction those cohorts do not publish.

PRIOR STRENGTH	A BAND CHANGES	B BAND CHANGES	C BAND CHANGES
K = 0	0	0	0
K = 2	0	0	0
K = 8	0	0	1

K = 0 removes shrinkage and reports Raw. The other alternatives test weaker and stronger pulls toward F. None changes a source, anchor or applicability mask. The results show how much of a band depends on the aggregation choice rather than on a new fact about a service.

Single-cell deletions test a different dependency. Across 55 deletions in A and 47 in B, the supplied outputs record 0 and 0 admission losses, and 0 and 0 eligibility losses. Their band-change counts are 0 and 0, subject to the small-cohort caveat above. In C, 17 of 68 deletions change a band, 0 lose row eligibility and 0 lose admission. The summary does not separately report loss of the cohort band gate; no such total is inferred here.

Deleting a low cell can raise a quotient. In the supplied C example, deleting Sightly's environment coverage cell, scored 1, raises its displayed VQ from 41 to 44. That is an

evidence-loss scenario, not an improvement in the service.

Missing-evidence scenarios instead fill a row's applicable n/e cells with 1, then 4, while holding F and every other row fixed. They change 0 bands in A, 0 in B and 6 in C. Imputation creates neither admission nor eligibility. These are scenarios, not bounds; the separate endpoint envelope uses 0 and 5. Lack of a band-change flag in an unbanded cohort is not evidence that missing cells have no numerical effect.

What the quotient is not

The VQ measures documented capability under the author's anchors. It supplies no estimate of accuracy, return on spend, legal safety or suitability for a particular contract. The authority ladder asks how far a verdict travels toward money. It does not imply that a reporting product has less practical value than a blocking product for every buyer. Legal allegations and incident records stay outside the index.

Claude agents scored the initial calibration rows with separate refute-mode checks. An OpenAI model scored the later rows, then self-checked them in a refute-mode pass; a supplemental consistency review also inspected those rows. The author checked the work and retained the primary scoring calls. There was no human second scorer, blind second score or inter-rater statistic. A second pass by a row's own scorer can correct a mistake while sharing the assumptions that caused it. Chapter 13 records that distinction.

The profiles

The profiles connect the purchased service to its entity, owner, history, method claims and sources. Read the profile with the matrix cell and the dated scope record before making a buying decision. A parent company's capabilities do not transfer automatically to a subsidiary's product, and a platform partnership cannot replace an unstated data route. The profiles also preserve contrary evidence and access limits, which a composite number cannot carry on its own.

Deep profiles — select a company

Each card shows the matrix scores as a bar strip. Open a company for its SWOT, what it sells, who buys it, how it earns, its verified capabilities with availability and evidence grade, performance evidence with sponsorship, limits, risks and fit. Pick a SWOT quadrant to isolate it and compare every profiled company on it. Each SWOT line is tagged record, vendor documentation, vendor claim or inference, and opportunities and threats are the author's reading of dated facts.

For each company, capabilities by dimension, performance evidence, financial record and timeline are in the web edition and in data/deep-profile-records.csv.

Adloox

OPERATING; ACQUIRED BY PEER39 (JUNE 2026, COMPLETION RECORDED IN A UK FILING); BRAND AND SERVICE RETAINED; MRC ACCREDITATION CONTINUED ON 29

JUNE 2026

P39 France SAS (RCS Paris 513 450 544; société par actions simplifiée à associé unique), formerly ADLOOX SAS (from 2011) and Scope3 France (June 2025 to July 2026), created as ADSAFE SARL with activity from 1 July 2009; UK sales subsidiary Adloox Limited (Companies House 08701942). MRC's letter of 29 June 2026 names the service 'Peer39 Adloox (formerly d/b/a Scope3 France Adloox) Viewability and Fraud prevention service' and is addressed to 'Peer39 P39 France'.

The suite run from Europe (MRC's letter is addressed to Paris): owned by Scope3 from November 2024 and by Peer39 since June 2026, with MRC accreditation continued under the new owner on 29 June 2026.

STRENGTHS

On 29 June 2026 MRC continued accreditation under Peer39: display and video impressions and viewability on desktop, mobile web and in-app, with SIVT on desktop and mobile web. RECORD [148] [63] [47]

Meta names Adloox for brand safety and suitability verification on Facebook and Instagram Feeds and Reels; Amazon lists it for viewability, brand safety and IVT on Amazon DSP. VENDOR DOCUMENTATION [358] [252]

Google's DV360 help documents Adloox pre-bid filtration and viewability settings, three on CTV; Adloox lists content-exclusion segments at $0.03 to $0.05 CPM and has targeted MFA since 2022. VENDOR DOCUMENTATION [359] [360] [361]

Among seeded cohort A candidates, only Adloox runs its MRC-accredited service from Europe; MRC addressed its June 2026 letter to P39 France in Paris. RECORD [148] [362]

TAG's registry lists Adloox with the Certified Against Fraud and Brand Safety Certified seals, both validated by a third-party auditor and covering global operations. RECORD [363]

WEAKNESSES

MRC accredits no Adloox pre-bid, CTV, brand-safety, Meta or Amazon metric, none is under review, and SIVT filtration is accredited on desktop and mobile web only. RECORD [47] [148] [65]

Accreditation claims exceed MRC's listed scope: Peer39's page implies accredited pre-bid viewability, and Adloox's 2024 releases announce unaccredited Meta and DV360 services under an 'MRC-accredited' descriptor, without disclaimers. VENDOR CLAIM [16] [364] [365] [47] [366]

Methods are summary-level, figures such as Peer39's IVT fall from about 12% to 3% lack denominators, and no independent labelled test of any Adloox category was found. VENDOR CLAIM [18] [367] [16] [368]

Peer39 says its post-buy reporting confirms 'that the pre-bid controls worked', controls Peer39 sells; Adloox also sells to publishers, and no separation policy was found. VENDOR DOCUMENTATION [16] [369] [370]

Adloox is documented on two of nine large platforms; YouTube, TikTok, Snap and Pinterest pages name other verifiers, and X, Reddit and LinkedIn were not screened. VENDOR DOCUMENTATION [371] [251] [372] [100] [261]

OPPORTUNITIES

Since June 2026 Peer39 has offered its customers Adloox verification in DV360, Amazon and Meta, a route to sell accredited post-serve measurement alongside its pre-bid contextual data. INFERENCE [25] [148] [373]

The Trade Desk's 2024 OpenPath terms deem MRC-accredited reports conclusive for invalid-

THREATS

Five months after Peer39's purchase, Google's DV360 help still lists 'Scope3 (formerly Adloox)' with a scope3.com contact, a naming lag that may confuse buyers and slow integration. INFERENCE [359] [378]

Adloox's Meta measurement runs on data Meta collects, filters and supplies, and access followed

inventory fees; Peer39 claims Adloox runs on The Trade Desk, where its accredited metrics could determine fees. INFERENCE [374] [25] [148]

In the 2023 French case Adloox declared it transfers no data outside the EU, and its terms place servers there, which could appeal to buyers requiring EU hosting. INFERENCE [375] [370]

Adloox has claimed OM SDK readiness for CTV since June 2023 and Google documents three Adloox CTV settings, so CTV is a candidate for new accredited metrics. INFERENCE [376] [359] [47]

AdCP's signal documentation names Peer39 as an example data provider and Peer39's 2026 release cites 'future-state agentic workflows'; Adloox verdicts could follow, though AdCP does not name Adloox. INFERENCE [377] [378]

refusals from 2016 to 2022; programme changes at Meta could reshape it. INFERENCE [375] [379] [358]

Peer39 describes DV360's 1 October 2026 change to inventory modes and content themes for first-party brand suitability; platform changes like this can alter what Adloox's segments control. INFERENCE [380] [359]

Peer39's page implying accredited pre-bid viewability could matter under MRC's 2024 policy on accreditation representations, which applies to the Peer39 Adloox service. INFERENCE [16] [366] [148]

Financial difficulty recorded by the French authority in 2023 and the UK subsidiary's net liabilities leave viability resting on an owner whose finances are unpublished. INFERENCE [375] [381]

WHAT IS SOLD

Adloox verification is sold as five lines, read on 10 November 2026. (1) Post-serve measurement by JavaScript, VAST and OM SDK tags: tracked ads, impressions, viewability and exposure time, GIVT and SIVT detection, and page-level brand safety and suitability classification, reported in a dashboard and as raw impression-level files (SFTP), with 'firewall' blocking of unwanted impressions. MRC accredits display and video tracked ads, impressions, viewability, exposure metrics and decision rate on desktop, mobile web and in-app, with SIVT filtration on desktop and mobile web only (in-app net of GIVT only); CTV measurement through OM SDK is claimed by Adloox and Peer39 but not accredited. (2) Pre-bid segments in DSPs: in Display & Video 360, documented by Google under the name 'Scope3 (formerly Adloox)': brand safety and suitability content filtration based on the GARM framework and seven content categories, fraud and invalid-traffic filtration (including MFA sites, per Adloox), both at URL, app, domain and subdomain level on open-web display and video, and viewability targeting by IAB viewed rate (no level stated), with three of the settings available on CTV; Adloox lists 36 DV360 segments, two groups of which carry list prices of $0.05 and $0.03 CPM; Peer39 also claims native pre-bid controls in Amazon DSP. None is MRC-accredited. (3) Walled-garden measurement: Meta (third-party brand safety and suitability verification on Facebook and

WHO BUYS IT

Advertisers, agencies, trading desks, publishers and ad platforms, mostly in France and Europe on the client stories read: advertisers Yves Rocher, Cofidis and SNCF Connect & Tech; agencies and trading desks Incubeta (UK), Numberly, Weborama, ZBO Media and NÜ Programmatique (Canada); publishers and ad networks Brut and Canal+ Brand Solutions; platforms Jun Group (US, using Adloox for its in-house DSP since early 2023) and Taboola's network (partnership of February 2023). Adloox's company page claims more than 400 brands live. ISBA's 2025 brand-safety guide (VER-ADLOOX-S66) names Adloox among example verification partners for UK advertisers. Since June 2026 the product is also offered to Peer39's customers, who 'can work with us now' for brand safety, viewability and fraud detection in DV360, Amazon and Meta.

Instagram Feeds and Reels, per Meta; viewability through Meta's viewability programme and an IVT rate, per Adloox; content blocklists documented by Meta under the Scope3 name (27 March 2025); Threads claimed by Peer39) and Amazon DSP (viewability, brand safety and fraud or IVT through server-to-server integration, per Amazon). (4) Publisher products: ad-server-activated exclusion segments and key-values for Google Ad Manager, an Equativ ad-server integration and publisher traffic measurement. (5) An attention solution (the CESP grid lists 'Adloox - Peer39') and MFA bulletins: attention belongs to the attention paper and is not scored here. Native placements are measured through Taboola's approved-vendor tag route. No audio or DOOH product is documented.

HOW IT EARNS

Fees per thousand impressions, on the documents read: DV360 pre-bid segment data fees set by the vendor and collected through Google ($0.05 CPM for brand suitability exclusion, $0.03 CPM for IAB unsafe categories, per Adloox's undated page); post-serve measurement billed by the vendor per served impression (Taboola's documentation, for its approved verification vendors generally). Peer39 presents its pre-bid controls as applied in DSPs (for example 2,000+ pre-bid categories in Amazon DSP, VER-ADLOOX-S73), advertises 'No Verification Fees' for the post-buy reporting that confirms those controls, and in a 2025 post promised page-level control across DSPs 'with no fees or minimums' (VER-ADLOOX-S51); how Peer39 charges for its segments, and how Adloox's separate measurement and segment fees fit that model after June 2026, are not stated on the pages read. No revenue, revenue mix or price list beyond these items is published; Adloox Limited (UK) files accounts without an income statement, and no French accounts were found in the registry data read. Neither acquisition price (Scope3, 2024; Peer39, 2026) was disclosed.

BEST FIT

European, especially French, advertisers and agencies wanting MRC-accredited post-serve viewability and desktop and mobile-web SIVT measurement with EU data hosting.

POOR FIT

Buyers needing accredited CTV measurement, accredited in-app SIVT, or any accredited brand-safety (property- or content-level) measurement.

Buyers needing third-party verification on YouTube, TikTok, Snap or Pinterest (not documented for

DV360 buyers wanting inexpensive pre-bid brand-suitability, fraud (including MFA) and viewability segments.

Buyers who want raw, impression-level verification files to join with their own logs (as Jun Group describes).

Peer39 customers adding post-serve verification and Meta and Amazon DSP coverage to pre-bid contextual controls.

Adloox).

Buyers who need published method documents, error rates or a written dispute and correction process before relying on a verdict.

Buyers wanting a verifier with no commercial stake in pre-bid data sales or publisher products.

Channel Factory ACTIVE; PRIVATE

Channel Factory, LLC (California limited liability company; SEC CIK 0001745403, Form D filed 9 July 2018). Its terms of 17 August 2026 and TAG's registry write 'Channel Factory LLC'; the privacy policy writes 'Channel Factory, LLC'.

A pre-bid YouTube classifier whose verdict is a list the platform executes: no MRC entry, an entity-level TAG seal, and a self-service plan (terms of 17 August 2026) that puts it in cohort C.

STRENGTHS

Google lists Channel Factory among four YouTube Activation Partners, a roster launched 6 October 2025 whose members, Google says, were reviewed and meet its standards. RECORD [382] [127] [383]

TAG's registry lists Channel Factory LLC as Brand Safety Certified, a CV Tool Vendor validated by a third-party auditor; the entity-level seal covers global operations. RECORD [384] [21]

kidSAFE has listed Channel Factory's ViewIQ Solution as a member since 2021 with +COPPA and SAFE certified seals, which concern child privacy and safety, not classification accuracy. RECORD [385]

Channel Factory says ViewIQ classifies 5 to 25 million videos a day with AI and human review, refreshes video-level decisions every 4 hours and covers 49+ languages. VENDOR CLAIM [19] [386] [387]

Channel Factory builds custom inclusion lists, including kid-friendly and 'Conscious' lists, and guarantees campaign KPIs such as CPV and VCR on buys it manages. VENDOR DOCUMENTATION [387] [388] [389] [390] [19]

WEAKNESSES

On MRC's accredited and under-review lists read 10 November 2026, Channel Factory, ViewIQ, ActivateIQ and AERO have no entry, so no sold unit is accredited or under review. RECORD [47] [65]

No independent error-rate test of ViewIQ was found; its 13+ categories and thresholds are unpublished, and its PwC-attributed result, worded three ways, links to a gated report. VENDOR CLAIM [19] [387] [386] [391]

ViewIQ verdicts reach the buy only as pre-bid lists; the August 2026 terms disclaim accuracy, tie no fee or credit to verdicts and offer no classification appeal. VENDOR DOCUMENTATION [19] [392] [23]

Google recognises Channel Factory for activation and targeting, not brand-safety reporting; DV360's YouTube measurement table and TikTok, Snap, Pinterest and LinkedIn pages name other brand-safety partners. RECORD [382] [393] [371] [394] [100] [261] [269]

Channel Factory builds suitability lists and runs managed media buying against them, offers creator branded content and sponsorships, and publishes no separation policy on the pages read. VENDOR DOCUMENTATION [19] [390] [389]

OPPORTUNITIES

Google's partner page, read 10 November 2026, says it is not onboarding new YouTube Activation Partners, leaving Channel Factory among only four on the roster. INFERENCE [127] [382] [383]

AI Slop Detection, launched 2 October 2026 with five pre-bid tiers at channel, video and creative level in 50+ languages, gives Channel Factory a new category to sell. INFERENCE [395]

The Premium CTV Intelligence Layer, announced 28 September 2026, could take Channel Factory's suitability classification beyond the large platforms; its publishers and devices are unnamed. INFERENCE [396] [397]

Channel Factory's August 2026 terms already bill a software platform monthly, and its 2023 site offered DV360 lists; pricing ViewIQ apart from media would separate lists from buying. INFERENCE [23] [392] [390]

TAG certification implies an auditor efficacy test and a wrongful-classification review process; publishing both would give buyers evidence and recourse that ViewIQ's public record lacks. INFERENCE [21] [384] [23]

THREATS

YouTube's derived-metrics policy of 14 September 2026 permits API-based safety and suitability assessment 'for the purpose of influencer marketing, vetting, and matchmaking'; whether ViewIQ's list-building qualifies is unstated. INFERENCE [398] [387]

Zefr, also on Google's activation roster, appears in YouTube's 2024 targeting and safety-reporting categories and DV360's brand-safety table, so buyers can get both roles from one vendor. INFERENCE [382] [393] [371]

A July 2024 House Judiciary staff report names Channel Factory as a GARM invitee and CISA 'common lexicon' collaborator; these are characterisations, not findings, but invite scrutiny. INFERENCE [283]

In mid-2023 Adalytics observed ads from iPhone 14 campaigns apparently run with Channel Factory on 'made for kids' channels; intent is unknown, but such audits test kids-content claims. INFERENCE [75] [387] [388]

Truelink invested in April 2025 without disclosing stake or control and said it would explore M&A; it also invested in the agency R/GA, a tie buyers may weigh. INFERENCE [40] [399]

WHAT IS SOLD

Product in the register: ViewIQ (aliases Channel Factory, Channel Factory ViewIQ; written 'View IQ' on the site), part of Channel Factory's 'IQ Series' with ActivateIQ and AERO. By product, metric, environment and placement, as documented on 10 November 2026: (1) ViewIQ, content suitability and brand safety classification on YouTube, pre-bid: Channel Factory says it uses YouTube's API and its ViewIQ engine to scan video, audio and metadata, classifies 5 to 25 million videos a day with AI plus human review, applies custom suitability filters in 13+ categories and advertiser-set risk thresholds, refreshes video-level decisions every 4 hours, scores 49+ languages, and delivers inclusion and exclusion lists of channels and videos (in 2023 up to 20,000 channels per category, with curated keyword block lists) that are executed as YouTube placement targeting (DV360 is named on the June 2023 site; the current pages name no buying tool).

WHO BUYS IT

Brands and agencies. Google describes its YouTube Activation Partners as serving 'both Advertisers and Agencies'. Named by Channel Factory: GroupM as a business customer supplying user data (privacy policy); logo image files named McDonald's, Tom Ford, Microsoft, Sony Music, Lyft and Starbucks (no market or period); case studies for McDonald's (Snapchat), McCafe (YouTube), ARAG, Vodafone, Greek National Tourism Organisation, Chester Zoo, OVO Energy and Elev8 clients; testimonials from William Grant & Sons (India, with DDB Mudramax) and Strawberry; a one-pager with Dentsu on a QSR brand; an analysis with OMG Sweden and Annalect covering 34 brands. Independent observation: Adalytics (2023) reports that Apple appeared to use Channel Factory for some iPhone 14 YouTube campaigns. Truelink's release cites 'deeply entrenched relationships with the global agency

Placement: YouTube videos and channels; device or environment scope is not stated. Includes kids-content avoidance or targeting and 'Conscious' inclusion lists (DE&I, sustainability, civic society). (2) Post-campaign reporting: a ViewIQ 'Inspect' screen (categories breakdown of a campaign's impressions; shown as a mock-up), ActivateIQ 'post-flight reporting', and a free 'content strategy analysis' of past campaigns from placement reports. No post-bid blocking, no impression-level tag, no IVT or viewability measurement. (3) ActivateIQ: AI-assisted media buying and optimization on YouTube, TikTok and other social platforms (video, audio and CTV formats), with 'Guaranteed Outcomes' on campaign KPIs (CPV, VCR, CPM). (4) AERO, 'Agentic Execution for ROI' (landing page undated; availability unconfirmed): brand-profile, list-building (YouTube and CTV), performance and reporting agents, claimed integrations with YouTube, Meta, TikTok, Snap and CTV. (5) Premium CTV Intelligence Layer (28 September 2026, IQ Series and AERO): video-level metadata tagging, genre and mood classification and contextual suitability activation across an unnamed premium publisher network, with reporting. (6) AI Slop Detection (2 October 2026): five-tier pre-bid classification at channel, video and creative level in 50+ languages, claimed across YouTube, Meta, TikTok, Reddit, Snapchat and CTV. Ways to buy: Self-Service Platform, Managed Services, Guaranteed Outcomes. Not sold: MRC-accredited measurement of any kind.

HOW IT EARNS

Not disclosed in figures: Channel Factory is private and its 2018 Form D marks its revenue range 'Decline to Disclose'. Documented modes: managed campaigns in which Channel Factory buys the media ('platform-native buying'), including 'Guaranteed Outcomes' with guaranteed CPMs (2023) or KPIs (2026); a self-service software platform whose terms (17 August 2026) provide free accounts and paid plans billed monthly in advance, non-refundable, with enterprise accounts invoiced and due within 30 days; a free placement analysis used for lead generation. No price list was found. Channel Factory markets NMSDC Diversity Media Credits (it shows an NMSDC 'Certified MBE 2020' badge), which lets clients count spend toward diversity goals. Inference, marked: on managed buys its revenue scales with the media bought holding companies'. No customer count or revenue split is published.

against its own lists; whether it charges a margin or a separate technology fee is not published.

BEST FIT

Advertisers buying YouTube through Channel Factory (managed or self-service) who want custom, values-based or kids-aware inclusion lists in many languages before the campaign runs.

Planners who want pre-bid suitability lists on YouTube and pair them with an MRC-accredited Brand Safety Reporting partner (Zefr for content-level YouTube reporting) for post-campaign verification.

Brands seeking NMSDC diversity media credits alongside suitability controls (a commercial attribute, not a verification capability).

POOR FIT

Buyers who need third-party, MRC-accredited verification of delivered impressions or a verdict that can support billing disputes, credits or make-goods.

Buyers who need open-web, non-platform in-app, audio or DOOH coverage, or invalid-traffic and viewability measurement.

Buyers who need the classifier to be separate from the party buying their media, or a published category taxonomy and error rates.

Buyers who need documented suitability controls on Meta, TikTok, Snap, Reddit, Pinterest, LinkedIn, Amazon or X.

DoubleVerify

ACTIVE, LISTED; MERGER WITH A SUBSIDIARY OF NEPTUNE BIDCO US INC. (NIELSEN'S PARENT) PENDING AT 10 NOVEMBER 2026

DoubleVerify Holdings, Inc. (Delaware; incorporated 16 August 2017; SEC registrant CIK 0001819928, Commission file 001-40349). MRC's accreditation letters are addressed to the operating company DoubleVerify, Inc.

The largest public verifier: $748.3 million of FY2025 revenue on about 9.5 trillion measured transactions, accredited for property-level verification in four environments; its sale to Neptune BidCo US Inc., Nielsen's parent, was agreed on 6 August 2026 and had not closed when read.

STRENGTHS

MRC accredits DV for SIVT-filtered impressions, pre-bid IVT detection and property-level ad verification on desktop, mobile web, in-app and CTV; on CTV no other service has property-level accreditation. RECORD [101] [47]

MRC accredits DV's viewability on desktop, mobile web, in-app and CTV video, and its third-party viewability reporting on Meta feeds, YouTube via Ads Data Hub and TikTok. RECORD [101] [47]

Meta, Google, TikTok, Amazon, Pinterest and LinkedIn each document DV as a third-party measurement partner on their own pages. RECORD [358] [243] [371] [400] [401] [261] [267]

WEAKNESSES

DV markets content-level controls on Meta, yet MRC accredits its suitability measurement only at property level, and no DV content-level service is accredited or under review. RECORD [403] [101] [47] [65] [404]

No independent error-rate test of a DV product was found; a 2026 study measured IAS-DV agreement at alpha 0.19, and DV's own accuracy figures carry no published method. RECORD [83] [176] [291] [191]

DV's product pages call the Authentic Ad 'a proprietary, MRC-accredited metric' covering intended geography; DV's own one-sheet says Authentic Ads are unaccredited when geographic

DV's FY2025 10-K reports $748.3 million of revenue, about 9.5 trillion Media Transactions Measured (its billing unit) and 131 customers each worth at least $1 million. RECORD [26]

DV publishes its main service agreement and partner-pay CPM list prices in a public Legal Center; the May 2026 list prices the Authentic Ad at $0.135 per thousand. VENDOR DOCUMENTATION [402]

targeting is included. VENDOR DOCUMENTATION [405] [406] [407] [404]

DV's Main Service Agreement waives fee disputes not raised within 30 days and provides data 'AS IS'; no verdict-dispute process or restatement policy was found. VENDOR DOCUMENTATION [402]

Supply-side revenue rose from 8.0% of revenue in FY2023 to 9.5% in FY2025, yet DV publishes no conflict or separation policy covering seller revenue or its prospective owner. RECORD [26] [408] [409] [133]

OPPORTUNITIES

The Trade Desk's OpenPath publisher terms make MRC-accredited third-party reports, a class including DV's, conclusive for Invalid Inventory fees; similar clauses elsewhere would give DV verdicts financial force. INFERENCE [374] [47]

A July 2026 DV blog describes an agency's AdCP buyer agent calling DV through MCP; a published transaction record would turn that vendor narrative into documented agentic use. INFERENCE [323] [334]

DV alone holds CTV property-level accreditation; the program-level signals Authentic Streaming TV added in January 2026 could, if audited, extend that record from app to program level. INFERENCE [47] [410]

MRC lists Meta's Third-Party Brand Suitability Pipeline as under review; if accredited, it could open a route to accrediting DV's Meta suitability measurement, now unaccredited. INFERENCE [65] [101] [358]

DV lists no audio or DOOH verification product; a 2023 audio fraud-scheme disclosure and June 2026 YouTube Audio Ads reporting give it a starting point in audio. INFERENCE [411] [412]

THREATS

Closing, now expected in Q4 2026, would put DV under Nielsen's parent with no published separation policy and would likely end SEC filings that disclose its supply-side revenue. INFERENCE [413] [133] [274] [26]

Adalytics' answer to DV's suit pleads truth, and a 2025 Senate letter asked the FTC to investigate verifiers including DV; neither is a finding, but both prolong scrutiny. INFERENCE [307] [312] [414]

Average Measured Transaction Fees fell 3% in 2025 (excluding one introductory deal) and 4% in the first nine months of 2026, pointing to price pressure as volume grew. INFERENCE [26] [317]

DV depends on partners: two programmatic platforms collected about 24% and 12% of 2025 revenue, and YouTube and Amazon control the data or environment DV measures. INFERENCE [26] [243] [101] [204]

MRC's list credits IAS, not DV, with S2S Amazon DSP accreditation (its chart disagrees), and DV360 offers only IAS brand-safety settings on CTV line items. RECORD [47] [61] [415] [359]

WHAT IS SOLD

One scored row, 'DoubleVerify advertiser verification' (cohort suite), sold inside the DV Media AdVantage Platform (DV MAP) and reported in DV Pinnacle. By product, metric, environment and placement, as documented on 10 November 2026: (1) DV Authentic Ad, the main sold metric, which the 10-K calls 'our definitive metric of digital media

WHO BUYS IT

Advertisers and their agencies are the main buyers: 90% of 2025 revenue (91% in 2024) came from advertiser customers, over 2,500 customers in total and 131 each worth at least $1 million of 2025 revenue (110 in 2024, 93 in 2023), no customer above 10%. Agencies 'have significant control over what solutions will be used' and some sell

quality', evaluating fraud, brand suitability, viewability and geography for each ad (partner-pay list price $0.135 CPM). (2) Fraud/IVT: post-bid GIVT and SIVT measurement and post-serve blocking on desktop, mobile web, in-app and CTV, and pre-bid avoidance segments and enterprise-level filtration delivered to DSPs and SSPs as files or APIs (all four environments MRC-accredited for SIVT filtration and pre-bid IVT detection); DV's pages list GIVT, SIVT, app and bot fraud, hijacked devices and injected ads. (3) Viewability: display viewability on desktop, mobile web and in-app; video viewability on those three and on CTV (certified apps only for viewable impressions); Fully On-Screen and quartile metrics on CTV; agency custom standards (GroupM, PMX); pre-bid viewability segments on desktop, mobile web and in-app and a Fully On-Screen segment on CTV. (4) Brand suitability: post-bid monitoring and blocking and pre-bid avoidance (Authentic Brand Suitability, a single customizable pre-bid segment; individual segments; DV360 tiers and 20+ categories) using DV Universal Content Intelligence and over 100 avoidance categories; MRC accredits it as property-level ad verification (176 languages binary, 56 domain, 41 page-level, 10 CTV), and DV's one-sheet says this is not content-level brand safety; MFA categories (since September 2023; high, medium and low tiers described in June 2024), 'Highly Illicit: Do Not Monetize' (February 2025) and Low Quality GenAI categories are sold but, per DV, not accredited. (5) Platform measurement (social and walled gardens): Meta (viewability on Facebook and Instagram feeds, accredited net of GIVT; content-level brand suitability measurement on Feeds and Reels since January 2024 and content-level controls since February 2025, which DV says let advertisers avoid content they deem unsuitable before their ads are served, unaccredited); YouTube (Ads Data Hub viewability for in-stream and bumper, accredited net of GIVT; brand safety and suitability including Shorts (December 2023) and Audio Ads (June 2026), unaccredited); TikTok (video viewability, GIVT and SIVT in-app, accredited April 2026; brand safety and suitability measurement and pre-bid video exclusion lists, unaccredited); X (viewability and Home Timeline adjacency in English- and Japanese-language markets, per a browser read the verifier pass could not repeat), Snap (viewability, attention, brand safety on three of five ad products, also per an unrepeated browser read), Pinterest (brand-safety reporting), LinkedIn

competing products. Two programmatic partner platforms collected about 24% and 12% of 2025 revenue on advertisers' behalf. Supply-side customers (publishers, retail media networks, ad networks and exchanges) generated 10% of 2025 revenue. Named users in vendor material include Butler/Till (agentic workflow blog) and Vodafone Germany and Lexus (case-study titles on the brand-suitability page); these are vendor case materials, not records of production scale.

(IVT, viewability, geography on feed and Audience Network; suitability on the Audience Network; CTV Ads), Amazon DSP (server-to-server verification) and Reddit (per DV's release), all unaccredited for DV on MRC's list. (6) CTV: DV Authentic Streaming TV (January 2026) adds program-level signals from licensed IMDb data (genre, maturity ratings, premiere dates) and replaces manual Do Not Air workflows with pre-bid enforcement through Authentic Brand Suitability; unaccredited. Adjacent products not in this row: Authentic Attention (the attention paper's), Scibids AI optimisation and DV Authentic AdVantage (MRC says Scibids is not audited or accredited), Rockerbox attribution (MTA, MMM, incrementality), DV Publisher Suite and DV Marketplace Suite (sell-side), DV Neura and DV AI Verification (agent-interaction reporting).

HOW IT EARNS

Advertisers pay a Measured Transaction Fee per thousand impressions measured on their behalf (the price of most solutions is fixed; on platforms that charge a percent of media, pricing is capped to mirror fixed fees); activation fees are often collected by DSP and platform partners and remitted to DV (partner-pay); supply-side customers pay monthly or annual contracts with minimum guarantees and tiered pricing. FY2025 revenue $748.3 million (Activation $427.3 million, Measurement $249.7 million, Supply-side $71.3 million); MTM rose 15% and the average fee fell 3% in 2025 (excluding an introductory fixed-fee deal); in the nine months to September 2026 revenue was $580.3 million, MTM +11% and fees -4%. DV publishes partner-pay CPM list prices (effective 21 May 2026): DV Authentic Ad $0.135, Authentic Ad Plus for Open Web $0.150 display and $0.370 video, Social Pre-Bid Controls $0.400, Video Complete $0.110, Authentic Attention for Open Web $0.050. DV's Main Service Agreement says 'DV usage statistics may be determinative for billing purposes' (its own fees) and waives fee disputes raised more than 30 days after the invoice date.

BEST FIT

Buyers who need MRC-accredited GIVT and SIVT filtration, viewability (on CTV, video only) and pre-bid IVT avoidance across desktop, mobile web, in-app and CTV through their DSPs.

POOR FIT

Buyers who require MRC-accredited content-level brand safety or suitability on any platform or the open web (DV holds none).

Buyers who need accredited third-party viewability on Meta feeds, YouTube in-stream (via ADH) or TikTok video, and documented partner integrations on the other large platforms.

Buyers who want property-level suitability with accredited multi-language coverage (176 languages for binary classification) and post-bid blocking on the open web and in apps.

Buyers who need a verdict-dispute or restatement process with time limits, or contractual remedies tied to DV's verdicts on the buy side (DV's MSA covers fee disputes only and disclaims accuracy).

Buyers who need published error rates or independently tested accuracy for a category.

Buyers needing accredited third-party measurement on X, Snap, Pinterest, LinkedIn, Reddit or Amazon, CTV display viewability, CTV pre-bid content suitability, audio or DOOH verification.

HUMAN ACTIVE; PRIVATE

HUMAN Security, Inc. (formerly White Ops), a Delaware corporation per its Authorized Third Party Data Use Agreement (updated 29 January 2025). MRC's letters are addressed to 'HUMAN'. No EDGAR registrant named HUMAN Security was found (10 November 2026).

The only cohort B candidate accredited for pre-bid and post-bid GIVT and SIVT on desktop, mobile web, in-app and CTV (April 2026); its signals sit inside platforms' own invalid-traffic filtering, and it has sold an advertiser suite since 16 June 2026.

STRENGTHS

On 14 April 2026 MRC continued HUMAN's GIVT and SIVT accreditation, pre-bid and post-bid, in four environments including CTV, and accredited viewability in desktop, mobile web and in-app. RECORD [150] [47]

MRC's pre-bid IVT detection table lists only HUMAN and DoubleVerify for CTV; IAS's row covers desktop, mobile web and in-app. RECORD [47]

Since March 2024 MRC's FreeWheel accreditation has included its HUMAN integration for up-front GIVT and SIVT filtration; DV360 and Amazon Ads use HUMAN inside their own IVT defenses. RECORD [416] [417] [197] [204]

HUMAN documents its GIVT and SIVT taxonomy and a closed loop: post-bid tag data retrains the pre-bid model every two hours, and a pass-through sample re-checks its predictions. VENDOR DOCUMENTATION [418] [419]

TAG's registry lists HUMAN as Certified Against Fraud through independent third-party validation, in the Fraud Detection Vendor category, for its global operations. RECORD [420]

WEAKNESSES

No source screened reports an independent miss rate, false-positive rate or precision for any category HUMAN sells; MRC says accredited vendors need not identify all SIVT. RECORD [82] [150] [421] [422]

Brand safety and suitability, on sale by February 2026, lacks MRC accreditation, a listed audit and TAG's Brand Safety Certified seal; its 'high accuracy' claim has no figures. RECORD [47] [65] [420] [421]

HUMAN's accreditation brief, still served in November 2026, claims accreditation for 'pre-bid mitigation of SIVT'; MRC accredits pre-bid IVT detection, not prevention of ad serving. RECORD [423] [82] [47]

MRC said in April 2025 that HUMAN does not directly filter requests, prevent ad serving or enable blocking; platforms decide, and payment follows their own determinations. RECORD [82] [197] [204]

HUMAN's terms disclaim accuracy, forbid sharing transaction-level reports and exclude 'Known Bad Actors' from even aggregate sharing; no process for disputing a verdict is published. VENDOR DOCUMENTATION [424] [425] [426] [427]

OPPORTUNITIES

The June 2026 suite offers HUMAN's IVT, viewability and brand safety to advertisers and agencies, beyond the DSPs and SSPs MRC named as pre-bid clients in 2025. INFERENCE [145] [82]

OpenX, first on the platform suite of 5 October 2026, uses HUMAN signals in supply curation and deal creation, carrying HUMAN's verdicts from bid filtering into packaged deals. INFERENCE [146]

OpenPath clause 4.4 and PubMatic section 8.4 let MRC-accredited vendors' IVT reports bear on fees or credits; HUMAN is accredited, though no such use is documented. INFERENCE [374] [228] [150]

Walmart Connect's Onsite Display audit, in process as of November 2025, includes SIVT performed by HUMAN; a grant would place HUMAN inside accredited retail-media measurement. INFERENCE [65]

THREATS

Senator Warner's March 2025 letters asked the FTC and Justice Department to investigate HUMAN, IAS and DoubleVerify over bot-detection claims; they are requests, not findings, but invite scrutiny. INFERENCE [132] [428] [429]

Adalytics' 2025 report observed ads served to declared bots by platforms citing HUMAN partnerships; it did not observe HUMAN's classifications, yet such failures can attach to HUMAN's name. INFERENCE [173] [132]

The Chavez suit, alleging HUMAN's bot-detection code collected visitor data on 1,200+ client sites, was dismissed with prejudice in July 2026 without a finding; similar claims could recur. INFERENCE [306] [308]

HUMAN sells to SSPs, exchanges and publishers as well as advertisers and publishes no conflict or separation policy, a gap that buyers weighing its verdicts may question. INFERENCE [430] [146] [145] [431]

HUMAN says much CTV inventory is delivered without JavaScript, leaving fewer signals, while schemes like NewsJunkie spoof devices at scale; CTV detection rests on thinner evidence. INFERENCE [432]

WHAT IS SOLD

HUMAN Ad Integrity Suite (sold to advertisers and agencies since 16 June 2026; Ad Integrity Suite for Platforms, for DSPs, SSPs and exchanges, since 5 October 2026 with OpenX first and further capabilities announced for H1 2027). By metric, method, environment and placement: (1) Invalid traffic, pre-bid: Ad Fraud Defense, documented as MediaGuard, an API or SDK that scores OpenRTB bid requests and returns IVT true or false with GIVT or SIVT taxonomy codes (plus optional, opt-in Policy flags such as made-for-advertising); the platform decides whether to bid; MRC-accredited for GIVT and SIVT requests in desktop, mobile web, in-app and CTV; MRC said in April 2025 these suggestions went to DSPs and SSPs, not advertisers. (2) Invalid traffic, post-bid: Ad Fraud Sensor, documented as FraudSensor, detection tags served with the impression (JavaScript tag, SDK or 1×1 pixel),

WHO BUYS IT

Trading platforms first: MRC said on 11 April 2025 that Ad Fraud Defense clients are generally DSPs and SSPs. Documented platform users: Google (DV360 integrates HUMAN as an extra safety check), Amazon Ads (HUMAN signals in its own IVT systems), LinkedIn (pre-bid filtering and post-bid detection from mid-2024), FreeWheel (MRC records its accredited measurement inclusive of HUMAN up-front filtration since March 2024), Magnite (Ad Fraud Defense since 2015), OpenX (platform suite, October 2026), Walmart Connect (SIVT by HUMAN in an MRC audit in process), Sovrn (viewability early adopter), AdRoll (quoted on the advertiser suite), and The Trade Desk (pre-bid SIVT blocking announced 31 August 2017; current use not confirmed by a Trade Desk page opened). Since 16 June 2026 the suite is also sold to advertisers and agencies; no advertiser client is named in the

reporting tracked ads, impressions, decisions, incomplete loads and GIVT and SIVT, MRC-accredited in the same four environments; MRC said in April 2025 HUMAN does not enable post-serve blocking. HUMAN's blog posts (2023 to 2025) also market Ad Fraud Defense and Ad Fraud Sensor for audio inventory (not accredited; no technical documentation). (3) Viewability: HUMAN Ad Viewability Measurement (FraudSensor-based, IVT-filtered), display and video in desktop, mobile web and in-app, MRC-accredited since 14 April 2026; not sold for CTV on the pages read. (4) Brand safety and suitability: AI classification of pages, apps or impressions with reason codes and URL-level visibility, a crawler (GlancelotBot) that retrieves page content, and custom brand-safety agents for platforms; HUMAN's brand-safety page (February 2026) also offers pre-bid targeting and blocking; no MRC accreditation or audit, no TAG Brand Safety Certified seal. (5) Made-for-advertising: domain and subdomain flags from combined indicators (as of August 2026), returned pre-bid as Policy flags. Reporting: dashboard with alerts, Reporting API and a log-level Snowflake data share. Outside this paper's metric families HUMAN also sells Ad Click Defense, Page Intelligence (landing-page IVT), Malvertising Defense and Ad Quality, and application-security products from the PerimeterX merger.

HOW IT EARNS

Subscription fees under order forms: the Subscription Agreement (updated 5 August 2025) invoices prepaid usage in advance and overage monthly in arrears, makes subscriptions non-cancellable and fees non-refundable, states that HUMAN's measurement of the client's usage is final for fees, and allows resale through authorised resellers; Malvertising Defense may be billed on 'Qualified Impressions'. No price list is published. Revenue is not filed: HUMAN is private and its parent's Form Ds (2024) mark the revenue range 'Decline to Disclose'; vendor statements say its revenue growth rate doubled in 2021 and the combined company would have more than $100 million in ARR (July 2022 merger announcement; all product lines). Payers include platforms using HUMAN on their own supply (sell-side and platform-paid), publishers, and since June 2026 advertisers releases opened. Customer counts are vendor statements: 500+ customers (combined, July 2022) and '500+ global brands' (October 2024), across all HUMAN product lines.

and agencies; revenue by customer type is not disclosed (H5: amount unresolved, not zero).

BEST FIT

DSPs, SSPs and exchanges that need MRC-accredited pre-bid GIVT and SIVT detection, including CTV, with an API and log-level data.

US buyers who want accredited post-bid IVT and IVT-filtered viewability measurement in desktop, mobile web and in-app from one tag, with reason codes.

Buyers and platforms that want IVT signals consistent with those already used inside Google, Amazon, LinkedIn or FreeWheel filtering.

POOR FIT

Buyers who need MRC-accredited content-level brand safety or suitability, or a published taxonomy and edition.

Buyers who need accredited CTV viewability, accredited IVT measurement in audio (HUMAN markets audio IVT protection without accreditation), or any documented DOOH IVT product.

Buyers who want vendor reports usable as evidence against counterparties without negotiation: the standard terms restrict sharing and disclaim accuracy.

Buyers who require a published, independent error rate for the detection they buy.

Integral Ad Science

ACTIVE; PRIVATELY HELD PORTFOLIO COMPANY OF NOVACAP-MANAGED FUNDS; NO LONGER FILES PERIODIC REPORTS

Integral Ad Science Holding Corp. (Delaware; SEC registrant CIK 0001842718 until its Form 15 of 2 January 2026). Services are contracted through Integral Ad Science, Inc. (Standard Terms); MRC's letters are addressed to 'Integral Ad Science'.

The other large full-suite verifier ($530.1 million of FY2024 revenue), with accredited integrations on Meta, YouTube and Amazon DSP; taken private by Novacap affiliates on 23 December 2025, so its last public figures end on 30 September 2025.

STRENGTHS

MRC's letter of 9 July 2026 continues IAS's viewability, SIVT and Quality Ads accreditation on desktop, mobile web and in-app, and CTV video metrics from certified applications. RECORD [67] [47]

In June 2026 MRC granted IAS initial accreditation of pre-bid viewability and IVT-avoidance segments in desktop, mobile web and in-app, and lists no other pre-bid viewability service. RECORD [63] [67] [47]

MRC accredits IAS's reporting on Meta (net of GIVT), YouTube via Ads Data Hub (GIVT+, not SIVT)

WEAKNESSES

MRC accredits no IAS content-level brand safety and no IAS measurement on TikTok, X, Snap, Pinterest, Reddit or LinkedIn; IAS's YouTube brand-safety audit has run since April 2023. RECORD [47] [65] [415]

No independent labelled error-rate test of IAS was found; IAS names no GARM edition, publishes no category definitions or thresholds, and its accuracy claims carry no method. RECORD [436] [126] [437] [27]

On Meta, MRC says IAS cannot independently identify IVT and has no user- or device-level data,

and, since November 2025, Amazon DSP server-to-server with GIVT and SIVT. RECORD [67] [66] [61] [47]

IAS measurement is documented on all nine large platforms, seven on the platforms' own pages; on TikTok, IAS's own tags measure viewability and IVT through the OM SDK. RECORD [358] [371] [433] [252] [99] [261] [268] [434] [263] [248]

TAG's registry lists IAS as Brand Safety Certified and Certified Against Fraud, both through independent third-party validation for global operations, and Certified for Transparency. RECORD [435]

though Total Net removes IVT matched by IAS's Fraud IP list. RECORD [67]

IAS's in-app page claims MRC-accredited viewability and SIVT without the letter's MRAID exclusion, and its ad-fraud page omits the letter's CTV certified-applications limit. RECORD [438] [439] [67]

IAS publishes no process to dispute a verdict or restate reports; its terms cap liability at 12 months' fees and exclude liability for inaccurate third-party data. VENDOR DOCUMENTATION [440] [441]

OPPORTUNITIES

Yahoo's display terms let an MRC-certified vendor's viewability report set advertisers' billable vCPM count, which Yahoo may dispute; IAS, certified with restrictions on Yahoo's list, could qualify. INFERENCE [223] [224]

The Trade Desk's OpenPath terms deem MRC-accredited reports conclusive in calculating publisher fees for invalid inventory; IAS's SIVT and viewability reports qualify, though no such use is documented. INFERENCE [374] [67]

A grant in MRC's YouTube brand-safety audit of IAS, open since April 2023, would give IAS the accredited content-level classification none of its services holds today. INFERENCE [65] [415]

IAS said on 27 April 2026 it would launch CTV pre-bid optimization; MRC's June 2026 pre-bid grants exclude CTV, which an audit of the new product could add. INFERENCE [442] [63] [67]

IAS announced brand-safety measurement for ChatGPT Ads in October 2026, after Threads, Pangle, X profiles, YouTube Audio Ads and Spotify podcasts; early entry could widen its platform reach. INFERENCE [355] [443] [444] [434] [445] [446]

THREATS

MRC's list shows IAS property-level ad verification that its July 2026 letter omits and its Q2 2026 report excludes; if MRC drops the row, IAS's listed accreditation narrows. INFERENCE [47] [67] [63]

Adalytics' 2024-2025 reports and 2025 Senate letters questioned IAS's brand-safety and bot detection; IAS answered each report, the letters are not findings, and new reports could renew scrutiny. INFERENCE [167] [169] [173] [447] [429] [168] [171] [178]

Publisher-channel revenue (14.4% in FY2024), partner-paid verification fees and an owner holding a digital media company, with no published separation policy, may lead buyers to question IAS's independence. INFERENCE [27] [448] [449] [450]

DSPs facilitated 46% of IAS's 2024 revenue and the largest DSP 18%, so a change in one DSP's integration or revenue share could cut IAS's revenue. INFERENCE [27]

A committed $1.0 billion term loan backed the take-private, and IAS has filed no periodic report since Q3 2025, so buyers cannot gauge its leverage. INFERENCE [277] [451] [452]

WHAT IS SOLD

Advertiser verification sold in two filed revenue channels (FY2024 10-K): Measurement (post-bid) and Optimization (pre-bid). Measurement: Firewall tags for verification, monitoring and, where enabled, blocking of display and video ads in desktop and mobile browsers, mobile in-app, CTV and digital audio, reported roughly 24 hours after delivery; Campaign Monitor pixel (browser); Account

WHO BUYS IT

Advertisers and agencies: over 2,000 advertising customers and 237 large customers (spending at least $200,000 a year) in 2024; 86% of FY2024 revenue came from advertiser customers. Served directly, through holding-company agencies (Dentsu, IPG, Omnicom, Publicis Groupe, WPP, Havas named in the 10-K) and through DSPs (Amazon, Display & Video 360, The Trade Desk,

Monitoring Tag; 'Proprietary Platforms' integrations sold as Total Media Quality (TMQ) for viewability, invalid traffic and brand safety and suitability on YouTube, TikTok, Meta and other platforms; Quality Impressions (viewable, valid, brand safe or suitable and in-geo per client settings); audibility; geo compliance; made-for-advertising and ad-clutter measurement (beta from 16 October 2023, later generally available). Optimization: pre-bid segments activated in DSPs (Amazon, Display & Video 360, The Trade Desk, Yahoo named in the 10-K) for viewability (40/50/60/70% segments), IVT avoidance (IP and user-agent segments, UA-only segments, Fraud IP list), brand safety and suitability (Context Control Avoidance, 100+ segments; Custom Brand Segments) and MFA avoidance; Quality Sync aligning pre-bid settings with post-bid measurement and blocking; Social Optimization inclusion and exclusion lists for YouTube, Meta and TikTok; YouTube Select lineups; Facebook brand-safety block lists (weekly); Meta Content Block Lists (test announced 10 October 2024); TikTok video exclusion lists; X pre-bid on Vertical Video only (US, opt-in). By platform and placement: Meta (Facebook and Instagram Feed, Stories, in-stream and Reels; Threads brand safety from 16 October 2025); YouTube (in-stream and bumper viewability via Ads Data Hub; brand safety incl. Shorts and, from 11 June 2026, audio ads, per IAS); TikTok (For You, Profile, Following and Search feeds, TikTok Lite, TikTok Ad Network/Pangle apps, CapCut, Lemon8); Amazon DSP (server-to-server, plus IAS JS and 1×1 tags on Amazon DSP creatives; Amazon O&O and third-party supply; Fire TV and Twitch outside the accredited scope); X (Home Timeline; profiles from January 2026); Snap; Pinterest; Reddit; LinkedIn (video IVT and viewability); Spotify (podcasts), Pandora (audio), Lyft Media (in-app), ChatGPT Ads (brand safety and suitability measurement announced 5 October 2026). CTV: video viewability, IVT and SSAI metrics (MRC: certified applications), TMQ for CTV brand safety and suitability via IRIS.TV video IDs, and IAS Total TV show, genre and rating transparency on Disney, NBCUniversal, Paramount, Prime Video and Publica publishers (27 April 2026). Outside this product row but sold by IAS: publisher tools (Publisher Optimization, Publisher Verification, Quality Connect in IAS Pulse), Publica CTV ad serving, and Quality Attention (the attention paper's subject); Lunio click-fraud protection is resold, delivered by Lunio. One row per product, metric, Yahoo), which facilitated 46% of 2024 revenue (largest DSP 18%). Publishers buy publisher tools and Publica CTV ad serving (Publisher channel, 14.4% of FY2024 revenue). A media partner can be designated to pay verification fees for an advertiser's campaign (Media Partner Fees). DSPs and other sell-side customers may resell IAS pre-bid segments. Geography: 69% Americas, 24% EMEA, 7% APAC (2024). Named users in vendor material: Mastercard with its agency dentsu (case study, 20 October 2023, no base or period).

environment and placement is in data table 'ias_sold_units_vs_mrc'.

HOW IT EARNS

CPM fees on volume: 'We recognize revenue by multiplying the cost per thousand impressions (CPM) and the number of impressions measured' (10-K). Pre-bid segment fees are invoiced through DSPs, with long-term revenue-share agreements under which DSPs including Google and The Trade Desk receive a percentage of revenue received through IAS products; Social Optimization is billed directly (YouTube on total tracked ads); IAS may invoice a third-party payor such as a media partner; fees are nonrefundable and IAS's usage statistics are determinative for its billing (Standard Terms). FY2024: revenue $530.1 million (Optimization $242.6m, Measurement $211.0m, Publisher $76.5m), net income $37.8 million, adjusted EBITDA (non-GAAP) $191.3 million; growth came from measurement impression volume (+23%) against average CPMs down 7%, and optimization volume +4% with CPMs +2%. Integration partners (platforms) produce no material revenue directly. Since 23 December 2025 IAS is private and has filed no FY2025 results; the last filed period is the nine months to 30 September 2025 (revenue $437.6 million).

BEST FIT

Advertisers buying programmatically across open web, in-app and CTV (certified apps) who want MRC-accredited viewability and SIVT, pre-bid viewability and IVT segments, and one vendor across DSPs.

Buyers on Amazon DSP, Meta and YouTube who want accredited third-party impression and viewability reporting inside those platforms' integrations.

Buyers needing post-bid brand-safety and suitability reporting on many social platforms (TikTok, Meta, X, Snap, Pinterest) under one GARM-aligned taxonomy.

POOR FIT

Buyers that require MRC-accredited content-level brand safety or suitability (none is held, on YouTube or elsewhere) or accredited property-level verification with a current letter.

Buyers that need published error rates, category definitions and thresholds, or a contractual route to dispute a verdict.

CTV buyers needing accredited pre-bid IVT or viewability outside certified applications, and audio or DOOH buyers needing accreditation.

Buyers who require published safeguards separating a verifier's buy-side verdicts from its sell-side businesses and its owner's media holdings.

Pixalate ACTIVE; PRIVATE

Pixalate, Inc., a Delaware corporation incorporated in 2011 (SEC Form D, CIK 0001582620, filed 29 July 2013). The Master Subscription Terms (12 February 2026) are offered by Pixalate, Inc. and MRC's letter of 8 April 2025 is addressed to it; the website Terms of Use (24 November 2025) are issued by 'Pixalate Europe Limited and its affiliates' under English law.

Accredited SIVT filtration on desktop, mobile web, in-app and CTV (letter of 8 April 2025), with pre-bid SIVT and video viewability listed in audit since September 2025.

STRENGTHS

MRC's 8 April 2025 letter continued accreditation of Pixalate's Analytics platform for impressions net of GIVT and SIVT, and SSAI video metrics, in four environments including CTV. RECORD [151] [60] [47]

Pixalate is one of five scored candidates whose SIVT detection and filtration MRC lists as accredited in CTV, with DoubleVerify, IAS, HUMAN and Protected Media. RECORD [47]

LinkedIn and MGID document Pixalate pre-bid checks on their own traffic, and Pixalate documents per-line-item GIVT and SIVT filtering in Amazon DSP since January 2022. VENDOR DOCUMENTATION [270] [453] [454] [455]

Pixalate publishes a coded 46-type IVT taxonomy, eight not yet live, its processing cadence, pre-bid probability bands, 13 versions of its subscription terms and a dated changelog. VENDOR DOCUMENTATION [456] [457] [458] [459] [460] [461]

Since 4 November 2025 Pixalate sells a Clawback Report: signed, hash-verifiable PDFs of accredited GIVT and SIVT ad counts by app or domain, seller, publisher and campaign. VENDOR DOCUMENTATION [462] [463]

WEAKNESSES

Pixalate's pre-bid SIVT feeds and Fraud API, in MRC audit since September 2025 with no decision, are unaccredited; MRC's pre-bid IVT table lists only DoubleVerify, HUMAN and IAS. RECORD [65] [64] [47]

Live Pixalate pages call pre-bid Blocking and MRT global-pool IVT rates 'MRC-accredited', beyond MRC's listed scope, and its SSAI metrics 'unique to Pixalate', though MRC lists IAS too. RECORD [455] [464] [465] [151] [47]

No independent labelled test of a Pixalate error rate was found; its benchmarks are own-classifier prevalence, and probability bands such as 'beyond reasonable doubt' lack published calibration. RECORD [466] [458] [467]

Pixalate's terms call IVT designations opinions and require its consent to use Analytics output in litigation or advertising claims, yet it markets Clawback Reports for refund negotiations. VENDOR DOCUMENTATION [459] [463] [462]

Pixalate reports no viewability in CTV, video viewability is a beta in MRC audit, brand safety is property-level and unaccredited, and no audio or DOOH product was found. VENDOR DOCUMENTATION [468] [469] [65] [464] [47] [470] [471]

OPPORTUNITIES

If MRC accredits pre-bid SIVT and video viewability, both in audit since September 2025, Pixalate would join DoubleVerify, HUMAN and IAS on MRC's pre-bid IVT list. INFERENCE [65] [47]

OpenPath and PubMatic terms give MRC-accredited IVT reports a role in fees or credits; Pixalate's

THREATS

Pixalate's live post calling pre-bid Blocking 'MRC-accredited' could draw scrutiny under MRC's 2024 policy, which authorises no accreditation claim for a service under review. INFERENCE [455] [366] [65]

Pixalate sells to SSPs and publishers while ranking sellers and labelling publishers; beyond narrow

signed Clawback Reports could serve there, though no such use is documented. INFERENCE [374] [228] [463] [151]

OpenEPG DB and its Pre-Bid API, launched in Q3 2026, add show-level CTV data beside accredited CTV SIVT, which could extend its CTV offer from apps to shows. INFERENCE [472] [151]

Pixalate's MCP connectors and IQ Agent let AI assistants query client IVT reports and ratings, a base for agent-driven buying, where no AdCP or AAMP implementation was found. INFERENCE [347] [470] [471]

Fraud API v3 has returned IVT reason codes since 13 October 2026, with Analytics, feeds and MRT to follow, which could make verdicts easier to explain and contest. INFERENCE [473] [456] [461]

safeguards, no separation policy was found, which buyers weighing its verdicts may question. INFERENCE [474] [475] [464] [476] [459]

Non-client sellers have an untimed, slower inquiry route, and Vidillion, an inventory seller, sued in 2018; the case ended at the pleading stage, and similar suits could recur. INFERENCE [477] [478] [479] [480] [481]

Pre-bid reach rests on platforms applying its lists with their own thresholds; Pixalate sued Criteo, a 2019 pre-bid partner, in April 2025, voluntarily dismissing in June. INFERENCE [482] [467] [483] [290]

Adalytics' 2024 report found Pixalate code in ads on MFA sites without knowing whether brands used its MFA tools; such case reports can attach failures to its name. INFERENCE [164] [484]

WHAT IS SOLD

(1) Pixalate Analytics, post-bid measurement (MRC-accredited as 'Pixalate's Analytics platform'): the client sends impression, click and video events through a JavaScript tag, a 1×1 pixel, a VAST impression pixel, a VAST verification script run through the OM SDK, a click tag or server-to-server; Pixalate reports GIVT and SIVT (gross, net of GIVT, total net of SIVT), ad counts and rendered impressions, SSAI video metrics, display viewability (JavaScript tag only), video viewability and engagement (beta; 54 metrics; VAST 3.0 or 4.x with OM SDK), click IVT, reach, MFA exposure, supply-chain (SCO) validation, privacy and COPPA reports and a mobile-app brand-safety report, in desktop web, mobile web, mobile in-app and CTV (CTV assigned by the device on which the tag fires; no viewability reported in CTV); formats: display, native (JS tag), video including SSAI. Since 4 November 2025 the Clawback Report issues signed PDFs of GIVT and SIVT ad counts by app or domain, seller, publisher, campaign and advertiser. (2) Pre-bid blocking: CSV data feeds the client applies on its own servers (IPv4 and IPv6 lists with IVT type and probability 0.5 to 1; data-centre ranges; mobile and CTV device IDs; High Risk Mobile and CTV Apps 2.0; High Risk Domains including a brand-safety list of known adult websites; delisted apps; malformed CTV bundle IDs; MFA lists; supply-path and user-agent lists), the Fraud API v3 (real-time probability plus IVT Taxonomy 2.0 risk-type codes for an IP,

WHO BUYS IT

Pixalate says its customers include agencies, DSPs, SSPs and publishers, and it does not disclose brand or agency clients without written approval. Its documentation addresses advertisers and agencies (Analytics, MRT, click fraud, MFA), publishers (Analytics on own traffic, MRT self-check, COPPA, Tag Server) and platforms (data feeds, SCO, COPPA and MFA screening, Fraud API). Platform users documented: Amazon DSP (Pixalate filtering selectable in line items since January 2022, per Pixalate), LinkedIn (Audience Network bid-request checks and IAB categories, per LinkedIn), MGID (pre-bid sampling and post-impression monitoring, per MGID), Criteo (global partnership announced 10 October 2019; Pixalate sued Criteo for breach of contract on 30 April 2025 and dismissed voluntarily on 18 June 2025), Verve (implementation announced 14 May 2019; declared interest); partner statements hosted by Pixalate add Pangle, the ad network of TikTok for Business (pre-bid technology and fraud monitoring, 19 May 2025), Screencore (Analytics, 22 October 2026), Gravite (pre-bid feeds and Analytics since March 2026, 5 November 2026) and, in 2024, zMaticoo, BIGO Ads, Broad Solutions, EscalaX and Anzu. OpenX is a lead only: Adalytics says OpenX's website lists Pixalate as a 'service provider' (not checked at OpenX). Yahoo lists Pixalate as a Certified Vendor on its Third Party Vendor List (5 November 2026). Enterprise services run under service orders and the Master

device ID or user agent), Pixalate filtering configured by Amazon in Amazon DSP line items, and a Tag Server that suppresses Google Ad Manager ad requests for publishers. (3) Media Ratings Terminal and Enrichment API: property-level ratings of mobile apps, CTV apps and websites (IVT from Pixalate's global pool, viewability estimates, brand-safety risk from app screenshots and text or root-domain analysis, MFA risk, COPPA audience and violation risk, authorized sellers, a Clawback Warning label since 6 October 2026), plus Seller and Publisher Trust Indexes; a free MRT tier exists. (4) Other: SCO Validation & Verification API, OpenEPG CTV show-level analytics (launched 16 June 2026) and enrichment (OpenEPG DB 1.0 and Pre-Bid API launched in Q3 2026; Enrichment API in beta), Analytics MCP, Enrichment MCP and IQ Agent for AI assistants. Cohort B scores the IVT measurement and filtering task; brand-safety ratings are property-level and outside cohort B's mask.

HOW IT EARNS

Subscriptions bought through service orders under the Master Subscription Terms (effective 12 February 2026; 13 versions published since before November 2021): fees invoiced and payable within 15 days, non-refundable, with a 10% annual increase on renewal, audits of client usage, incremental charges for unannounced 'Events Spikes', and Analytics Reports API use capped at 1 TB processed a month with $0.10 per GB above it. Pre-bid data feeds are marketed as an on-premise solution 'at a Flat-Rate'. Ad Trust and Safety API plans under the website Terms of Use run in billing cycles (for example monthly, quarterly or annual) with overage fees and no refunds on cancellation (Terms of Use). Client data feed Pixalate's pooled detection by default; siloed data is available but 'changes pricing'. No price list, revenue figure or revenue split by customer type is published; the 2013 Form D marks the revenue range 'Decline to Disclose'.

BEST FIT

DSPs, SSPs and exchanges that want pre-bid IVT lists or a real-time API across CTV, mobile apps and web, with per-entry probabilities and their own thresholds.

POOR FIT

Buyers who need MRC-accredited pre-bid IVT protection, CTV viewability or content-level brand safety.

Buyers who need published, independent error rates for IVT detection.

Buyers and sellers who need MRC-accredited post-bid GIVT and SIVT measurement and SSAI metrics across desktop, mobile web, in-app and CTV from one tag or server-to-server feed.

Supply-side teams that need documented IVT evidence for clawback negotiations, within the limits of Pixalate's terms.

Buyers vetting mobile and CTV apps for IVT, COPPA, MFA and supply-path risk before buying (MRT, Enrichment API).

Audio or DOOH buyers (no product found).

Parties who need to use verification output in litigation or advertising claims without the vendor's consent.

Sellers who need a timed, formal appeal against an IVT designation.

Zefr ACTIVE, PRIVATE

ZEFR, Inc., a Delaware corporation (SEC EDGAR registrant 'ZEFR, INC.', CIK 0001474764, which records the former name MOVIECLIPS, INC. for filings from 3 November 2009 to 16 August 2012). Zefr's website terms name 'ZEFR, Inc.'; kidSAFE's member page lists the publisher as 'Zefr, Inc.'. MRC's letters are addressed to Rich Raddon, Co-Founder and Co-CEO, Zefr.

The only third party with MRC content-level brand safety and suitability accreditation, for YouTube alone (English-language in-stream and music video; letter of 20 January 2026); it sells the same classification, unaccredited, on Meta, TikTok, Snap and Reddit.

STRENGTHS

On MRC's list of 10 November 2026, Zefr is the only third party accredited for content-level brand safety and suitability, covering English-language YouTube in-stream video and music video. RECORD [47] [152]

Zefr's content-level grant of 20 January 2026 covers desktop, mobile web, in-app and CTV; no other content-level service on MRC's list is accredited on CTV. RECORD [152] [47]

On 6 November 2025 MRC accredited Zefr's processing of Google Ads Data Hub YouTube video impressions and viewability, net of GIVT, on desktop, mobile web and in-app. RECORD [153] [61] [47]

Google, Meta, TikTok and Snap name Zefr for brand safety or suitability measurement on their own pages, and TikTok and Meta let it manage pre-bid exclusion lists. RECORD [371] [243] [358] [485] [251] [486] [487] [100]

Zefr acquired Adverif.ai, a misinformation-detection company, in July 2022, and has since sold a misinformation category on TikTok (February 2023)

WEAKNESSES

MRC accredits Zefr's content-level service for English-language YouTube in-stream video alone; its Shorts, Google Video Partners, Meta, TikTok, Snap and Reddit products are neither accredited nor under review. RECORD [152] [47] [65] [415] [490] [491] [492]

Zefr claims 'MRC accreditation for viewability & IVT' and IVT coverage on TikTok; MRC accredits YouTube viewability net of GIVT only, and no TikTok page documents Zefr IVT. RECORD [493] [494] [153] [47] [251] [486] [487]

No independent error-rate test of a Zefr product was found, and Zefr publishes no category definitions, thresholds, error statistics or process for disputing a verdict. RECORD [83] [368] [283] [495] [496] [497] [498]

Zefr measures YouTube campaigns while buying YouTube media for advertisers, and has supplied signals to TikTok's own Inventory Filter; it publishes no separation policy or revenue mix. VENDOR DOCUMENTATION [499] [500] [493] [501] [502]

and YouTube (September 2024). VENDOR CLAIM [488] [489] [490]

After GARM's August 2024 discontinuation, Zefr's later 2024 releases call its twelve categories industry-aligned without naming GARM; its Media Services page still says 'GARM-aligned', naming no edition. VENDOR DOCUMENTATION [503] [490] [504] [499] [496]

OPPORTUNITIES

MRC's 2018 supplement requires direct measurement or validated platform data for third-party content-level accreditation; Meta's partner data pipeline, under MRC review, could supply it for Zefr on Meta. INFERENCE [3] [65]

MRC's October 2025 policy asks verifiers to submit content-level products for audit; Zefr's Shorts, Google Video Partners and non-English YouTube reporting, already sold, could widen its grant. INFERENCE [11] [152] [490] [491] [495]

Zefr's current solution and technology pages do not mention CTV, though MRC accredits its YouTube content-level service there; marketing that scope could reach YouTube CTV buyers. INFERENCE [152] [47] [37] [495]

Zefr's Zain hub, in closed alpha per its September 2026 release, has an AdCP-based YouTube buying integration and an MCP server; a named transaction would document agentic use. INFERENCE [343] [344]

If Reddit documents the partnership Zefr announced on 30 September 2026, Zefr would be named by five large platforms; no Reddit page read on 10 November names it. INFERENCE [492] [505]

THREATS

IAS's YouTube brand-safety service, covering English and Spanish content including CTV, has been in MRC audit since April 2023; a grant would give Zefr an accredited competitor. INFERENCE [65] [415]

Zefr's YouTube reporting runs on Google's Ads Data Hub, and TikTok and Meta execute its exclusion lists; a change to any platform's partner programme could reshape Zefr's products. INFERENCE [153] [244] [371] [486] [485]

DoubleVerify and IAS hold matching brand-safety roles on YouTube, Meta, TikTok and Snap, so Zefr's platform positions are not exclusive; Pinterest, Amazon, LinkedIn and X name other verifiers. INFERENCE [371] [245] [485] [251] [486] [100] [261] [204] [270] [506]

MRC's October 2025 policy withdraws accreditation from vendors that label text- or keyword-based property-level verification Brand Safety; its reach over Zefr's domain-level 'Pre-Screen Brand Safety' product is unresolved. INFERENCE [11] [507] [508]

A July 2024 House Judiciary staff report called Zefr a GARM partner and disputed one of its classifications; it is not a finding, but similar inquiries could recur. INFERENCE [283]

WHAT IS SOLD

By product, metric, environment and placement, as documented on 10 November 2026. (1) Brand Safety & Suitability (sold as Zefr Data Solutions and inside Zefr Media Services): content-level classification of the videos and posts next to which an advertiser's ads ran or may run, against Zefr's suitability categories (GARM-mapped by name in its releases through 2023), sold as post-campaign measurement and reporting in the Atrium dashboard and as pre-bid or pre-impression exclusion. YouTube: post-campaign reporting through Google Ads Data Hub, MRC-accredited for English-

WHO BUYS IT

Advertisers and agencies, by Zefr's account and platform listings. Named in vendor materials: Mazda with its US media agency Garage Team Mazda (YouTube, 2021 managed-service targeting campaign), a pharmaceutical brand with CMI Media Group (Meta; case study published 14 September 2023), a pharmaceutical brand on YouTube (2025), Waitrose and M&S Food (home-page testimonials), Rexona (TikTok's own case tile); IPG Mediabrands is quoted in Zefr's September 2024 YouTube misinformation release and Omnicom Media Group executives in 2022 and 2023 releases. Platforms list

language in-stream video and music video content on desktop, mobile web, in-app and CTV; Zefr also sells reporting for Shorts (documented from September 2024), Google Video Partners (general availability 29 January 2026) and YouTube Audio Ads (11 June 2026), which the accreditation does not list; inclusion and exclusion lists for Google Ads and DV360 campaigns. Meta: third-party brand suitability verification on Facebook Feed (available 30 March 2023), Instagram Feed (Zefr's repost of 2 August 2023; Meta's page lists an update dated 31 July 2023) and Facebook and Instagram Reels (10 October 2023), and third-party content block lists (13 February 2025, per Meta); Zefr adds custom publisher (placement-level) exclusion lists, including for Audience Network (Zefr, 10 October 2024), pre-impression avoidance (its 'Zefr Max' product, February 2025) and content block lists for Threads feed, which it measured first (May 2026, per Zefr only). TikTok: post-bid measurement since January 2022 (a badged TikTok Marketing Partner per Zefr's release of 23 February 2023 and TikTok's announcement of 19 May 2023), partner-managed Video Exclusion Lists that TikTok calls the pre-bid control (Standard and Limited inventory), and, per Zefr (21 April 2026), coverage of Search, Smart+, TikTok Lite and GMV Max including viewability and 'invalid traffic (IVT)'. Snap: post-bid brand safety and suitability measurement (Zefr, 29 October 2024; Snap names Zefr, 24 February 2025); Snap's Approved Vendors page lists Zefr for Brand Safety on Snap Ads Video, Commercials and Story Ads, not Lenses or Sponsored Snaps, and not for viewability. Reddit: partnership announced by Zefr on 30 September 2026 (no Reddit page opened names Zefr). Google Search Partner Network: weekly domain-level pre-screen exclusion lists built from the domain list Google shares (Google Ads Help; Zefr, 8 May 2025). (2) Viewability: YouTube Ads Data Hub video impressions and viewability for skippable and non-skippable in-stream and bumper ads, MRC-accredited on desktop, mobile web and in-app net of GIVT only (Shorts, in-feed and masthead excluded); Zefr also offers viewability 'across social platforms' at 'no additional cost' (a vendor claim: platform pages document Zefr viewability only on YouTube through Ads Data Hub). (3) Brand Misappropriation, AI-Generated Content and Media Quality modules (vendor pages only). (4) Zefr Media Services: managed YouTube buying with 'GARM-aligned brand safety enforcement', priced as Zefr for advertisers to enable (Google DV360 and Google Ads, Meta, TikTok, Snap). Regional resellers: Mitsui and its subsidiary Legoliss in Japan (April 2021); Black C Media (Türkiye) and RedC Media (GCC and Egypt) from July 2026. ISBA's UK Responsible Media Guide (V1, March 2025) names Zefr among example ad verification partners. Zefr claims 'Fortune 100 Brands' as users (vendor statement).

one fee. (5) Zefr Labs custom AI; Opti bid optimisation for YouTube Demand Gen and Performance Max (6 August 2026); Zain, an agentic hub with an 'AdCP-based' YouTube buying integration and an MCP server, in closed alpha (14 September 2026).

HOW IT EARNS

Not disclosed in any filing: Zefr is private and its Form Ds 'Decline to Disclose' a revenue range. Zefr said in April 2021 that it was 'a profitable technology company ... with over $100MM in annual revenue' (vendor statement; period and basis not stated). Its pages say viewability is provided at 'no additional cost' and 'free of charge', and that Media Services is priced as 'One Simple Fee' covering the managed service; no CPM, fee schedule or contract is published. Google's DV360 help warns that enabling a measurement vendor may bring 'additional charges' from the vendor. Zefr's 2020 partnership with IRIS.TV, a sell-side contextual marketplace, was reported by Adweek (reposted by Zefr) to include a revenue-sharing agreement when their data is used for contextual buying; whether it continues is not stated. Who pays for the free viewability metric, and whether any platform pays Zefr, is not stated in the sources opened.

BEST FIT

Advertisers needing third-party, content-level suitability reporting on YouTube in-stream inventory in English, including CTV, with an MRC-accredited basis [VER-ZEFR-S02].

Buyers on TikTok and Meta who want partner-managed, video-level or content-level exclusion lists executed by the platform before the bid [VER-ZEFR-S73; VER-ZEFR-S66].

Brands in sensitive categories (pharmaceutical, kids, alcohol) wanting custom suitability profiles on social video, accepting vendor-only evidence of accuracy [VER-ZEFR-S20; VER-ZEFR-S63].

POOR FIT

Buyers who need accredited measurement outside YouTube, or on YouTube Shorts, Google Video Partners or non-English content [VER-ZEFR-S01; VER-ZEFR-S02].

Buyers who need SIVT or fraud measurement, or viewability on CTV [VER-ZEFR-S03].

Open-web display or CTV app buying outside YouTube (only a domain-level Search Partner Network product was found) [VER-ZEFR-S70].

Buyers who need published error rates, method documents or a formal dispute process before treating a verdict as settlement evidence [VER-ZEFR-S24; VER-ZEFR-S28].

CHAPTER 12

What a buyer should ask for

Before relying on a verification verdict, ask for three separate records for the exact service bought. An accreditation letter establishes its audited scope, a labelled test establishes an error estimate, and the contract states any deciding role in payment. The screened record to 10 November 2026 supplied no independent error rate for a scored product. Two of sixteen readable platform and exchange terms gave a third-party report determining effect, within limited contractual scopes. A buyer's executed contract may differ.

> **IN SHORT**
>
> Media Rating Council (MRC) accreditation establishes standards conformance within scope; the independent studies screened supplied no sold-category error rate for a scored product ◆ ◆. Two readable terms gave third-party reports determining effect, with no settlement found ◆ ◆. Platform data routes and accredited connected-TV (CTV) metrics require separate checks ◆ ◆. Seller revenue and ownership call for disclosed safeguards; the DoubleVerify separation statement found is limited ◆ ◆. Product pages naming GARM identify no dated implementation ◆ ◆. Agent-workflow reporting still leaves the verdict-to-decision link and financial agreement unestablished ◆ [137, 138].

Eight questions, chapter by chapter

Statuses were read by 10 November 2026 and will be rechecked at the freeze on 11 December 2026.

First, ask for MRC's letter for the exact unit bought: product, metric, environment, format, placement, language and standard (Chapter 2). Several public terms that give a third-party report a role in payment, among them The Trade Desk's OpenPath clause and PubMatic's fraud credit, ask for an MRC-accredited service ◆ ◆. MRC names the letter as the authoritative reference, and the list does not always match: IAS's property-level row has no letter on MRC's site ◆ ◆. Letters also narrow scope: Zefr's content-level accreditation covers English-language in-stream video and music video on YouTube, not Shorts ◆. Peer39's verification page implies accredited pre-bid viewability that MRC does not list ◆.

Second, request the miss or false-positive rate for the bought category, the labelled test set, settings, period and denominator (Chapter 3). The independent screen found cases and agreement studies without a qualifying error rate ◆. Protected-versus-unprotected comparisons measured by the protecting vendor cannot supply that test on their own ◆. Ask who assigned the labels.

Third, read who decides, pays and triggers a fee adjustment (Chapter 4). OpenPath makes accredited reports conclusive for The Trade Desk's publisher payments, while Microsoft announced the end of Invest's partner-report refund route ◆ ◆. PubMatic retains a good-faith decision on fraud credits ◆. H2 is supported on its registered test, the listed platform terms; the wider contract population is reported as sensitivity (Chapter 4) ◆. IAB's May 2026 terms leave controlling measurement to the order ◆.

Fourth, on a large platform, ask whose data the verdict is made from (Chapter 5). A verifier given only the platform's counts cannot detect invalid traffic on devices it never sees, a limit MRC records for IAS on Meta ◆. Google accepts no third-party pixels for YouTube measurement, and Amazon says the verifiers' detection runs inside its environment ◆ ◆. On Meta and YouTube, MRC accredits the verifiers' viewability net of general invalid traffic (GIVT) only and the platforms' own measurement including SIVT ◆ ◆ ◆. TikTok is the exception: DoubleVerify and IAS said in 2021 and 2022 that their tag ran there through a measurement kit TikTok controls; no 2026 page restates it ◆.

Fifth, on CTV, ask whether MRC accredits the metric bought there, and for what share of the buy (Chapter 6). CTV SIVT filtration is accredited for five scored products, but CTV video viewability only for DoubleVerify and IAS, whose viewable impressions count certified apps only, and neither letter says what share of the vendor's CTV traffic those apps carry ◆ ◆ ◆. Content-level accreditation on CTV is Zefr's alone, for YouTube, and the program-level signals vendors sell there carry none ◆ ◆. MRC's 2021 guidance does not require detection of ads played while the TV is off, though it asks vendors to disclose the limit ◆.

Sixth, ask who else pays the verifier, and for the written policy that keeps those payers apart from the verdict (Chapter 7). Sellers brought DoubleVerify 8.6% and IAS 14.4% of FY2024 revenue, IAS's line including Publica's CTV ad serving ◆. IAS's terms let a media partner pay an advertiser's verification fees, so a seller can pay for the check on its own inventory ◆. DoubleVerify's separation page of June 2024 names no scope, owner or enforcement and does not address its agreed sale to Neptune BidCo US Inc., Nielsen's parent, which had no closing on record on 10 November 2026 ◆ ◆. The author works at Samba TV, which competes with Nielsen in TV measurement. Seller revenue and ownership are governance attributes, not evidence of bias ◆.

Seventh, request the implemented framework, dated edition and maintainer (Chapter 8). The WFA records June 2022 as GARM's latest update, while the product pages screened name no implemented edition ◆ ◆ ◆. MRC's content-level supplement lets buyers choose risk tolerance and defines no tiers of its own ◆.

Last, ask which protocol state is documented and which agreement gives the verdict financial force (Chapter 10). AdCP's thresholds can name the verifier, with the author's working-group interest declared ◆. DoubleVerify and Zefr make production claims; new reporting partly corroborates Butler/Till's quality-analysis workflow without proving that a verdict drove the agent's purchase ◆ ◆ [137, 138]. The screened record establishes no such financial-authority agreement ◆.

The decision tool pairs each question with the evidence that would answer it.

Before a verdict settles anything: what to ask for

Pick the question a buyer is asking. The panel shows the evidence that would answer it, what that evidence assumes, where it fails and what it still cannot tell you.

BUYER QUESTION	TEST THAT FITS	WHAT THE INFERENCE NEEDS	WHERE IT FAILS	WHAT IT CANNOT TELL YOU
Is the verdict I am buying accredited for the unit I buy?	MRC's accredited and under-review lists, and the letter itself, for the exact product, metric, environment, format, placement and language bought, read on the day of purchase.	That the letter is current and its exclusions (placements, languages, pre-bid versus post-bid) are read, not only the company's name on the list.	Where the product sells a unit MRC does not list: content-level classification outside YouTube, viewability where only impressions and invalid traffic are accredited, pre-bid segments still in audit, platform placements a letter excludes.	Whether the verdict is right. Accreditation certifies conformance of a scoped measurement process; the public letter does not provide a labelled-set product error rate.
How often is the product wrong in the category I buy?	An independent test of that product with a labelled set and a denominator, for the category and setting bought, published with the product's settings and the observation period.	That the labels are independent of the vendors tested and that the sample resembles the inventory bought.	In the sources screened to 10 November 2026 no such test of a scored product's sold category was found. Case studies, prevalence in a sample and agreement between vendors measure other things.	Agreement between two vendors is not accuracy, and a vendor's own invalid-traffic rate is not its error rate.
Does the verdict decide what I pay?	The contract or platform term that gives a named third-party report, or a class of them, determining effect on the billable quantity, a fee, a credit or a remedy, with who decides, who pays and what triggers it.	That the executed contract matches the public term, and that the platform does not keep the final say.	The large platforms bill and credit on their own measurement. On the list fixed at registration, two readable terms gave a third-party report determining effect, one of them in a buying route since retired, and the other binds a platform and its publishers rather than advertisers.	Whether a settlement was ever paid on a third-party verdict: none was found in the sources screened, and executed contracts are private.
Inside a large platform, whose data is the verdict made from?	The platform's and the verifier's documents naming the data route for each metric: data the platform supplies, a vendor tag the platform fires, or the verifier's own code in the placement.	That the documents are current; platform pages change often.	On eight of the nine large platforms the documented route is data the platform supplies, or no route is documented. The one verifier-executed route, on TikTok, rests on vendor documents from 2021 and 2022.	Whether the platform's data is complete: a verifier reading a platform's feed cannot see what the feed leaves out.
Who else pays my verifier, and what keeps that apart?	The verifier's filed revenue by customer type, its owner's other interests, and any published conflict or separation policy that covers them.	That a filed revenue category says who pays; it does not say who pays for each verdict.	Private verifiers file nothing, and IAS stopped filing after it was taken private in December 2025.	Whether seller revenue or ownership changes any verdict. Revenue and ownership are governance facts, not evidence of bias.
Which suitability framework does the product implement?	The dated edition the product maps to, who publishes it, and how it changes.	That "GARM-aligned" is read as alignment to the archived June 2022 edition, which no longer has a maintainer.	GARM was discontinued in August 2024. Later documents reproduce its floor or publish their own guidance; which of them, if any, maintains floor and tiers with a change process is set out in chapter 8.	Whether a product's classification is right on a given page or video.
Can an agent act on the verdict in a buy?	The protocol state documented for that verdict (schema field, connector, interoperability test, production use, demonstrated transaction) and any agreement that gives it financial force.	That membership and schema fields are not read as deployment.	The production claims found are vendors' own statements, with no named campaign record or transaction.	Whether the verdict is accurate, or binding on anyone.

What would settle the open questions

- An independent, published test of a scored product's error rate on a sold category, with a labelled test set, a denominator and the vendor's response ◆.
- A settlement, credit or make-good paid on a third-party verdict, or a term that gives such a verdict determining effect on what an advertiser pays ◆ ◆.

- MRC accreditation of content-level brand safety or suitability for a scored product on a large platform other than YouTube, on the open web or on CTV outside YouTube; Meta's pipeline that supplies suitability data to third parties is under MRC review ◆.
- A named body publishing a dated floor-and-tiers suitability framework after June 2022 with a change process; the one candidate found, a guide by ISBA, the UK advertisers' body, whose brand-safety chapter is dated March 2025, pledges review but publishes no procedure (Chapter 8) ◆ ◆.
- A verification verdict in documented production use in an agent-driven buy, with an agreement that gives it financial force ◆.

CHAPTER 13

Method

The paper was registered before it was researched: its hypotheses, scoring rules and anchors were fixed on 9 November 2026 after an outside AI review of the brief whose factual points were checked at their primary sources. Claude agents gathered the workstreams and profiles and scored the calibration rows with separate refute-mode verifiers. An OpenAI model completed the later sweep checks, scored and self-checked the added rows, and fact-checked the chapters. The author checked the work and made every inclusion, cohort and verdict call. Sources were read on 9–11 November 2026; status-sensitive facts are due for recheck at the planned freeze on 11 December. The process supplies no human second score or audit.

IN SHORT

The brief went through the author's own fact check on 2 November 2026 and an outside review by an OpenAI model on 4 November (144 entries); the hypotheses and scoring rules were fixed on 9 November, before any research agent ran. Claude research agents produced 10 workstreams, a vendor-universe sweep, 7 company profiles and 7 analyses of strengths, weaknesses, opportunities and threats (SWOTs). Separate Claude verifiers, told to refute, checked the workstreams, initial sweep and deep profiles; an OpenAI model checked the later sweep files. Across the verifier records, of 1125 checks, 696 found the item supported as stated, 420 supported it with a qualification, 9 contradicted it and 0 left it unresolved. The research registered 1458 source records, 59 of them readable only in part, and 804 claims. Products are scored on 12 dimensions: Claude agents scored the calibration rows with separate refute-mode checks; an OpenAI model scored the added rows and then self-checked them in refute mode, and the Verdict Quotient (VQ) summarises documented capability within a cohort in bands, never ranks. Executed contracts, log-level data and vendors' test sets are private, and no vendor or third party was contacted ◆ ◆. The author works at Samba, a Nielsen competitor in TV measurement, and co-leads a working group of AgenticAdvertising.org, which maintains the Ad Context Protocol (AdCP).

Fixed before the research

Brief v1 was drafted on 2 November 2026 from the author's published record. The same day, nine checker agents in the author's AI-assisted process, not independent or external

reviewers, each opened the primary sources behind one group of its dated facts, and v2 corrected what they found.

On 4 November 2026 the author ran an outside review of v2 by an OpenAI model, which returned 144 entries. Before the brief changed, nine more checker agents re-opened the primary sources behind 87 of them: 71 were confirmed, 15 were right only in part and one was wrong, a challenge to the date White Ops renamed itself HUMAN, 30 March 2021, which the original release confirms ◆. The checkers also caught errors the reviewer had missed: "100% brand safe" is the phrase of Adalytics' August 2024 report ◆, and the letter the Media Rating Council (MRC) now links for the property-level row of Integral Ad Science (IAS) does not cover ad verification ◆.

Brief v3, dated 9 November 2026, records each disposition and changed the method as well as the facts. The hypotheses gained numeric decision rules: H1 was split into accredited coverage and claims against scope, H2 into four components with first-party validation recorded apart, and H3 was coded by the estimand each study supports. All twelve scoring dimensions were rewritten so that each attribute has one home, with anchors for every value from 0 to 5, and the legal and incident record moved outside the quotient. Three of the reviewer's scoring proposals were adopted in a different form: a credible failed test still scores below marketing claims, though a 0 now needs a labelled test with a denominator; the authority dimension keeps its ladder toward settlement, as the paper's question; and non-disclosure scores 1, while ownership earns or loses nothing on its own.

The hypothesis register and the scoring rules were fixed on 9 November 2026, before any research agent ran or any cell was opened. The register holds 12 entries: hypotheses H1a, H1b and H2 to H7, each with its population, comparator, observable outcome and rules for supported, contradicted and unresolved, and questions Q1 to Q4, each with a coding rule. Absence of accessible evidence yields unresolved, never contradicted. The product register, seeded from the brief and completed by the sweep, holds 45 records, 19 of them listed and 9 historical.

How the evidence was gathered

Claude research agents gathered the original evidence under the author's direction on 9 and 10 November 2026. The OpenAI run-1 work on 10–11 November added sweep checks, scoring and chapter fact-checks; the author checked the combined work. Working from the shared rules, the brief and the registers, each wrote one structured file: sources with dates, funding and access limits; findings with the supporting passage, evidence class, confidence, limitations and counterevidence; and data tables whose rows name their sources. The rules require every cited source to be opened, treat a search snippet as a lead, and settle each claim from the source suited to its type, such as filings for ownership and the accreditor's record for accreditation. Where a page refused automated fetches, agents retried with a browser user agent or in a browser and recorded the route; X's help pages were read in a browser ◆ ◆.

10 workstreams, A to J, each took one part of the brief, and 12 vendor-universe sweep, U, named every company and product it could find.

STREAM	WHAT IT READ	BEARS ON
A	MRC's lists, letters and standards; vendors' accreditation claims; certification and framework bodies	H1a, H1b, H6, Q4
B	The Association of National Advertisers' (ANA's) 2023 study, Adalytics' reports and the vendors' answers, academic work, platform transparency reports	H3
C	Vendor and commissioned claims of detection, accuracy and savings, with sponsor, design, sample and period	H3, kept apart from B
D	Filings, measured volume, ownership, funding and consolidation since 2021	H5, Q3
E	Billing, credit and remedy terms on H2's fixed list, and public contracts	H2
F	What third-party verification can see and do inside the nine large platforms	H4
G	What is measured and accredited on connected TV (CTV) and in-app	Q2, Q4, H1a; research question 10
H	Letters, court filings, regulator actions and independent reports since 2023	Q1
I	What buyers, agencies and their bodies document about verification use, fees, contracts and disputes	H2
J	Verification verdicts in AdCP, the Agentic Advertising Management Protocol (AAMP) and vendor agent interfaces	H7
U	Every named company and product, for the product register	The cohorts

7 deep profiles covered the companies the brief chose for market weight and depth of public evidence: DoubleVerify, IAS, HUMAN, Pixalate, Zefr, Channel Factory and Adloox, which Peer39 bought from Scope3 in June 2026 ◆. Claude agents produced those profiles with separate refute-mode verifiers. Each records the entity, its products, buyers, revenue model, financials and timeline, and a capability record per scoring dimension; 7 SWOTs summarise them, each point tied to its sources and labelled as a record, vendor documentation, a vendor claim or inference. A second-wave sweep opened the remaining candidate products' documentation. An OpenAI model completed the later sweep verifications and scored and self-checked the added rows; its chapter fact-checks and the combined work were checked by the author.

The research registered 1458 source records, including those the verifiers added: 798 vendor pages, documents and releases, 98 filings and 36 studies among them. 59 could be read only in part or as an abstract, and claims resting on them are limited to what could be read. Each finding became one claim in the ledger, with one evidence class; a vendor's documentation is a vendor assertion however technical, a study of unknown funding is not treated as independent, and the author's essays never count as corroboration.

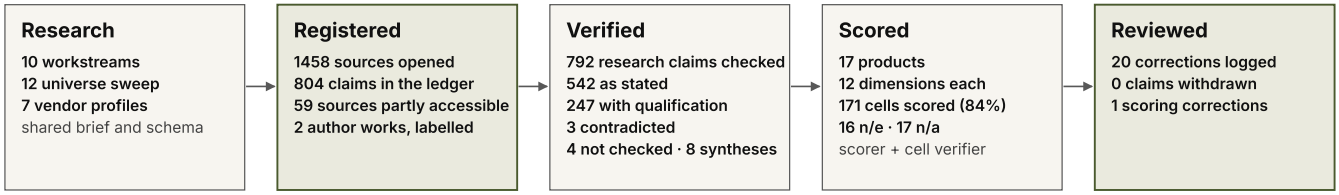
EVIDENCE CLASS	WHAT IT COVERS	CLAIMS
Direct record	A filing, registry entry, court record or other official document	272
Independent measurement	Measurement by a party with no financial interest	26
Commissioned measurement	Measurement paid for by an interested party	10
Vendor assertion	The vendor's own pages, documentation and releases	255
Synthesis	A reading that combines sources	209
Inference	A reasoned conclusion, marked as such	28
Forecast	A party's statement about the future	4

How it was checked

In the verifier pass, the original Claude workstream files, universe sweep and deep profiles went to separate Claude verifiers told to refute them. An OpenAI model completed the later

capability-sweep verifications. These file checks differ from the added score rows, whose scorers performed their own refute-mode second pass. The verifier re-opened the cited sources, checked each number's denominator, unit, date, entity and version, and returned one verdict per item, with the correction, the source's wording or the reason wherever the item fell short. Verifiers made 1125 checks of findings, profile capability records, profile fields and table rows: 696 supported as stated, 420 with a qualification, 9 contradicted and 0 unresolved. The figure traces each stage.

THE PIPELINE, AND WHAT EACH STAGE PRODUCED OR REMOVED



What the process is not: two agents reading one source are not two measurements; a model checking a model is not an audit; a cell verifier re-reading the scorer's sources is not a human inter-rater test.

Every count on this figure is computed from the ledgers at build time (data/stats.json); the prose quotes the same values.

Research cutoff 2026-11-11. Claims outside the verifier sample are marked unresolved with a note; they are not counted as supported.

> **Every stage was built to remove things, and the counts are computed, not typed.** The five stages of this project with the actual outputs of each. The verifier sample covers the material claims per file, not every claim; the rest are labelled unresolved rather than assumed. Corrections include withdrawn claims and score changes from the review passes.
>
> Units: counts · Scope: this research project, as of the cutoff
>
> Counts describe process volume, not evidence quality; a claim that survived verification is supported by its source, not proven true.

Each claim is printed in its final wording, the verifier's correction where there is one. Of the 804 claims, 792 were checked: 550 stand as stated, 247 with a qualification and 3 were contradicted, which are cited only to say so. The 4 unresolved claims rest partly on HUMAN press releases that refused the verifier's fetches; the parts it re-opened held, and the rest is cited only as unresolved.

The contradicted claims show what the pass is for. The verifier contradicted stream C's top of the vendor-reported invalid-traffic range, 39%, with a 46.1% row in the stream's own table ◆. It contradicted stream I's reading that MRC's 2018 content-level supplement left error disclosure optional, since the supplement makes several such disclosures mandatory ◆. It also contradicted the sweep's proposal to list Pixability as bundled with its media, having found a self-service model sold apart ◆. The verifiers had limits too: one had used up its web-search budget and re-bounded absences only through pages it opened directly, and another could not repeat a profile's browser reads of X and Snap ◆. A model checking a model is not an audit, and two agents reading one source are not two measurements.

Every claim marked in the text is listed below with its evidence class, review outcome, confidence and supporting passages.

The claim ledger lists the 426 claims marked in the text. The print edition leaves it out; it is listed in full in the web edition and in data/claim-ledger.csv, with each claim's evidence class, review outcome and confidence.

How the products are scored

Each row is a named product serving one buying task. Company-wide evidence does not automatically transfer between products, and aliases create no extra row. Cohort A measures viewability, general and sophisticated invalid traffic (GIVT and SIVT), and safety or suitability across at least two environments. Cohort B sells impression IVT measurement or filtration; cohort C sells content safety or suitability for verification or blocking. Platform first-party measurement, researchers and standards bodies are described separately ◆.

The author made every inclusion, cohort and verdict call. The product register records each inclusion reason, and the corrections ledger logs every change after the verifier pass, starting with the register decisions of 10 November 2026, taken before any cell was scored. The sweep had proposed two rows for HUMAN's Ad Integrity Suite, which would have brought cohort A to the size that bands need; its verifier objected that the brief forbids splitting a product to reach it, and HUMAN describes one suite ◆ [145, 146]. The suite is scored once, in cohort B, where its primary output and accredited scope sit. Channel Factory's ViewIQ entered cohort C on terms of 17 August 2026 that sell it apart from media [23]; Pixability and Sightly entered on the same test, pending the second wave's check of their buying task [22, 20].

The dimensions, anchors and masks were fixed on 9 November 2026: cohort A takes all twelve dimensions, B ten and C eleven, and n/a (not applicable) comes only from the mask.

GROUP AND WEIGHT	DIMENSION	WHAT IT ASKS	COHORTS
Depth, 0.45	Method disclosure	How each sold metric is determined, as published	A, B, C
Depth, 0.45	Independent detection evidence	Published error-rate tests with a labelled test set and a denominator	A, B, C
Depth, 0.45	Invalid-traffic method	GIVT and SIVT scope, pre-bid and post-bid, per environment sold	A, B
Depth, 0.45	Safety and suitability method	Framework and dated edition, content-level or placement-level, languages and media	A, C
Depth, 0.45	Verdict uses and authority	How far the verdict travels toward settlement, from report only to a documented settlement	A, B, C
Reach, 0.25	Environment coverage	Live environments among desktop, mobile web, in-app, CTV, audio and digital out-of-home	A, B, C
Reach, 0.25	Platform access	Documented partnership on the nine large platforms	A, C
Reach, 0.25	Measured scale	Impressions or transactions measured in a stated 12-month period	A, B, C
Reach, 0.25	US availability and use	Documented US availability and production use	A, B, C
Standing, 0.30	Accreditation	Accredited units as a share of sold units, on MRC's records	A, B, C
Standing, 0.30	Interest disclosure and safeguards	Revenue by customer type, sell-side products, a conflict or separation policy	A, B, C
Standing, 0.30	Dispute and correction process	A dispute route with time limits, a correction policy, documented corrections	A, B, C

A 0 requires the anchor's documented absence or failure, never an allegation. Missing evidence is n/e. Cells record availability, date, rationale, sources and evidence grade; vendor technical documentation is capped at 4 and marketing at 3. The masks alone determine n/a. Announced or retired products are listed, and admission needs four counted cells.

Claude agents scored the calibration rows with separate refute-mode verifiers. An OpenAI model scored the added run-1 rows and then self-checked each in a refute-mode pass. A supplemental consistency review inspected the added rows and led to the Protected dispute-process correction. The author checked the work and retained the primary scores. Only verified and qualified cells count; an unverified cell is treated as n/e. A scorer's self-check is not a blind second score. There is no human second scorer and no inter-rater statistic.

Cross-row interpretations after extraction

Cross-row interpretations identified as R1 to R8 were made after extraction on 10 November 2026. They are post-extraction interpretations, not rules fixed before research: R1, an environment counts as sold where the vendor says it measures or protects advertisers' traffic there; R2, off-platform ad networks such as Pangle and Audience Network are in-app inventory, not platforms; R3, one H1a unit per product, metric family and environment, pre-bid and post-bid counted once; R4, OpenPath clause 4.4 gives anchor 4 only where the product is MRC-accredited for the relevant metric and documented as utilized by The Trade Desk or its Ad Partners; record the evidence grade for both conditions and apply the applicable cap; R5, a vendor's own volume statement with a unit and period scores at grade D; R6, sources a scorer opens are registered before they are cited; R7, one standard for every row, whatever the vendor's size; R8, Pixability and Sightly are scored only if sold apart from managed media, which the sweep verifiers confirmed. The original registered anchors remain distinguishable from those later interpretations.

The Verdict Quotient

The VQ combines documented capability within a cohort. It averages counted cells within Depth, Reach and Standing, skipping n/e, n/a and unverified cells. Raw is their weighted mean at 0.45/0.25/0.30, renormalised if a group is missing. F is the mean Raw score of admitted cohort rows, including thin records. Shrunk = (k × Raw + K × F) / (k + K), with counted-cell total k and prior strength K = 4. VQ = 100 × Shrunk / 5, with halves rounded up only for display. Accuracy and legal safety remain outside the index.

K is the author's heuristic rather than a fitted statistical parameter. Shrinkage moderates thin records' observed averages without imputing missing cells. The VQ borrows the Contextual Quotient's band offsets but uses a different field mean; no validation follows from their resemblance.

Bands come from the unrounded offset 20 × (Shrunk − F), and equal values share a band.

BAND	OFFSET, 20 × (SHRUNK − F)
Most documented	+12 or more
Well documented	+1 or more, below +12

BAND	OFFSET, 20 × (SHRUNK − F)
Partly documented	−8 or more, below +1
Least documented	Below −8

Band eligibility requires 6 counted cells and one in each applicable group. A cohort needs 6 eligible rows for bands; rows missing a group show which one, without a band. Chapter 11 reports the tests below.

TEST	WHAT IT CHANGES	WHAT IS REPORTED
Reweighting	Depth, Reach and Standing at equal thirds, 0.60/0.20/0.20, 0.30/0.45/0.25 and 0.30/0.20/0.50, with Raw and F recomputed across the fixed cohort	How far quotients and bands move
Prior strength	K = 0 (Raw), 2 and 8	How far quotients and bands move
One-cell deletion	Each counted cell turned into n/e in turn, with k and eligibility recomputed	Quotient change, admission loss and eligibility loss, separately
Missing-evidence scenarios	Every applicable n/e cell of one row filled with 1, then 4, with F and every other row frozen	Rows whose band changes; imputed cells never create admission or eligibility
Envelope	The same fills at the scale's ends, 0 and 5	Reported apart, never used to flag a band
Coverage	Counted cells as a share of applicable dimensions	Reported for every row

A band changed by one-cell deletion or a missing-evidence scenario is marked sensitive. Such ordering supports no buyer recommendation. A dated adverse-record 0 remains supported by its archive unless superseded.

The review of the finished paper

The finished paper was submitted to an outside AI review. Review findings and any resulting changes must be recorded in the corrections ledger before publication.

The corrections ledger holds every substantive change made after the verifier pass, with what questioned it, what was checked, the verdict, the change and how it was verified.

> The corrections ledger has 20 entries. The print edition leaves it out; it is listed in full in the web edition and in data/corrections-ledger.csv, with the original wording, the challenge, the evidence checked, the verdict, the change made, how it was implemented and how it was verified.

Limits

Public records leave executed contracts, recovered credits, logs and vendor test sets largely inaccessible. The vendor/commissioned screen found no sold-category error rate on a labelled holdout with a denominator ◆; the pricing screen found no matched comparison of verified and unverified impressions ◆. Advertisers supplied logs to some studies, including the ANA's 2023 study ◆. Sparse public disclosure can understate a product's capability.

Status-sensitive facts were read on 9 and 10 November 2026 and are dated in the text: MRC's list carries no update stamp, and DoubleVerify's sale was pending at the last check ◆

◆. The evidence cutoff is 11 November 2026; those facts are read again on 11 December 2026, and later changes are to be logged separately.

Absences are bounded to the logged source screen and date. Failed access establishes no absence. The legal record excludes sealed and nonpublic matters and does not cover every state docket (Chapter 9).

Agents misread sources. Separate refute-mode verifiers checked the original Claude research files; an OpenAI model completed later sweep checks and chapter fact-checks and self-checked its added scoring rows. The author checked the combined work. These checks can share model or source assumptions. Corrections are logged, scores have no human second scorer, and the US dimension counts documented US availability only.

Interests

The author has worked since May 2025 at Samba, the TV measurement and advertising-data company formerly branded Samba TV. Nielsen competes with Samba in TV measurement, and on 6 August 2026 DoubleVerify agreed to merge with a subsidiary of Neptune BidCo US Inc., Nielsen's parent ◆; the merger was pending at the last check, on 10 November 2026 ◆. The paper says so wherever the transaction is discussed at length, and DoubleVerify is scored under the same published rules as every other vendor.

He co-leads the Signals & Measurement working group of AgenticAdvertising.org, which maintains AdCP and whose scope touches verification signals, so workstream J's findings on AdCP are a declared interest. No Fluff Advisory, his firm, operates a signals agent and a sales agent in that organisation's registry [6]; in records read on 10 November 2026, the sales agent is one of two that passed a storyboard testing billing after a post-SIVT window, which names a fictional measurement vendor ◆. Comments he filed on IAB Tech Lab's OpenProposal, the AAMP 3.0 draft ◆, are registered as his own work, an interest rather than evidence, and none concerns verification fields [139].

Earlier he was an executive vice president at Verve Group, an advertising platform whose supply the verifiers measure, from September 2022 to August 2024; Verve announced in May 2019 that it had implemented Pixalate's invalid-traffic detection, and no source opened shows whether that continued ◆. He has also worked at Hearts & Science, an agency that bought verification for its clients, and at Lotame. His essays arguing that a verification verdict "moves nothing" and that advertising has no adjudicator for disputed conversion counts are rival framings tested by H2, not evidence.

No vendor funded, reviewed or saw the paper, and no vendor or third party was contacted while it was prepared, so no correction invitation could be mistaken for editorial control. The paper, its data and its review prompts name no company as a consulting client. Every vendor is scored under the same published rules.

APPENDIX

Appendix: hypotheses, products and downloads

The registered hypotheses and how they read

The hypotheses and their decision rules were registered on 9 November 2026. The readings below apply to November evidence, before the planned 11 December 2026 freeze recheck. The author's 11 November rulings retain H1b, H5, H6 and H7's chapter readings and settle H2 on its listed-terms rule. Their alternative interpretations remain part of the record.

ID	HYPOTHESIS	READING	BASIS	CHAPTER
H1a	At least three quarters of sold units hold MRC accreditation	Contradicted	108/298 pooled (36.2%); A 80/179 (44.7%), B 25/82 (30.5%), C 3/37 (8.1%); unknown denominators remain unknown ◆	2
H1b	Product accreditation claims stay within listed scope	Contradicted, author's interpretation	Peer39's pre-bid wording and Protected's CTV display claim exceed scope on the retained readings; no MRC finding is inferred ◆	2
H2	Third-party verdicts rarely determine payment	Supported on the listed-terms rule	2/16 readable entries (12.5%); wider-contract and qualified-clause sensitivity retained; no qualifying paid settlement or matched price difference found in the screen ◆	4
H3	No independent public sold-category error rate is established for a scored product	Supported within the screen	No qualifying labelled-set test with a published denominator found; IAS's publisher-tool candidate is outside the scored task ◆	3
H4	Every documented platform route uses platform-supplied data	Contradicted, medium confidence	TikTok's dated measurement-kit route meets the verifier-executed comparator; platform control and route freshness qualify it ◆	5
H5	Seller revenue or interested ownership has no published safeguard	Contradicted, narrowly	Author accepts DoubleVerify's separation statement for publisher clients; stricter policy reading retained ◆	7
H6	No maintained successor floor-and-tiers framework exists	Contradicted on the author's process reading	ISBA's pledge, overseer and trigger accepted; formal-procedure reading unresolved; product implementation clause unestablished ◆	8
H7	Verification verdicts stop at schemas, without production use or financial authority	Unresolved	Connectors and tests exceed schema-only use; production evidence lacks the verdict-to-decision link and an identified financial agreement ◆	10

The coded questions

ID	QUESTION	WHERE ANSWERED
Q1	Which letters, proceedings, filings and reports concern the products, and what did vendors answer?	Chapter 9, outside the quotient
Q2	What does each product measure, through which method, at which stage, and with what action and coverage?	Chapter 1, with platform routes in Chapter 5 and CTV and in-app detail in Chapter 6
Q3	How has the market consolidated since 2021?	Chapter 7, keeping agreements, approvals, completions, rebrands and retirements distinct
Q4	Which standards and versions govern the services and assurance?	Chapter 2, with CTV and in-app specifications in Chapter 6

The product universe

The register records each company and purchased service, its owner, cohort, status and inclusion reason. Listed and historical records are described without receiving a zero score. The matrix compares products within their registered buying task. Mantis Sentinel is listed, not scored: the final review found neither separate purchase terms nor a live or beta status for it.

Product register: scored, listed and historical

Every product the vendor-universe sweep named, with the cohort decision of 10 November 2026. Inclusion reasons are in data/vendor-universe.csv and data/product-register.csv.

Machine-readable copy: data/vendor-universe.csv.

PRODUCT	OWNER	OPERATING ENTITY	PEER GROUP	WHAT IT SELLS	IN THIS PAPER	US AVAILABILITY	STATUS AT CUTOFF	SRC
DoubleVerify advertiser verification	Public (NYSE: DV) at 10 Nov 2026	DoubleVerify Holdings, Inc. and subsidiaries	Verification: third-party verification companies	Pre-bid avoidance and post-bid measurement and blocking of invalid traffic, brand suitability and viewability for advertisers and agencies (buying task: full verification suite)	Scored in cohort A	yes (VER-U-S81: New York dateline; SEC registrant per VER-U-S42)	live	[7] [47] [799] [413] [754]
IAS advertiser verification	Novacap (completion announced 23 Dec 2025; IAS stock ceased trading on Nasdaq)	Integral Ad Science, Inc.	Verification: third-party verification companies	Pre-bid filtering, post-bid blocking and measurement of invalid traffic, brand safety and suitability and viewability for advertisers (buying task: full verification suite)	Scored in cohort A	yes (VER-U-S44: New York dateline)	live	[8] [47] [65] [800]
Adloox verification	Peer39 (acquired from Scope3; Peer39 blog of 3 Jun 2026); Peer39 self-described as 'Wholly independent, employee owned'...	Peer39 P39 France (MRC letter addressee); Adloox	Verification: third-party verification companies	Post-bid measurement of invalid traffic and viewability with brand-safety controls, plus pre- and post-bid controls in DV360, Amazon DSP and Meta, for advertisers, agencies and buy-side technologies...	Scored in cohort A	yes (VER-U-S31: Cambridge, MA address; VER-U-S35)	live	[516] [148] [47] [25] [760] [521]
Protected by Mediaocean	Mediaocean LLC: Flashtalking acquired Protected Media in 2021; Mediaocean agreed to acquire Flashtalking on 13 Jul 2021...	Protected Media (Protected by Mediaocean)	Verification: third-party verification companies	Viewability, invalid-traffic and property-level brand-safety measurement and blocking for advertisers and agencies, and supply protection for publishers and platforms (buying task: verification suite)	Scored in cohort A	yes (VER-U-S06: letter addressed to 120 Broadway, New York)	live	[30] [635] [149] [50] [824] [47]

PRODUCT	OWNER	OPERATING ENTITY	PEER GROUP	WHAT IT SELLS	IN THIS PAPER	US AVAILABILITY	STATUS AT CUTOFF	SRC
mFilterIt Ad Traffic Validation	Unconfirmed (no ownership statement on the pages opened, 10 Nov 2026)	mFilterIt (US entity: MFILTERIT TECHNOLOGIES INC, Northborough, Massachusetts)	Verification: third-party verification companies	Pre-bid avoidance and post-bid measurement of GIVT and SIVT, viewability and attention measurement, MFA detection, and brand safety and suitability with post-bid placement-level verification, for...	Scored in cohort A	unconfirmed (VER-U-S90: US entity and address listed; no US production use...	live	[32] [766] [767] [47] [65]
HUMAN Ad Integrity Suite	Unconfirmed: private; current control not stated in any source opened (10 Nov 2026)	HUMAN Security, Inc.	Fraud: invalid-traffic and bot-security companies	Pre-bid scoring of OpenRTB bid requests and post-bid detection tags that filter GIVT and SIVT, for DSPs, SSPs, exchanges and buyers (buying task: IVT filtration on the buy side's behalf)	Scored in cohort B	yes (VER-U-S03: letter addressed to 111 W. 33rd Street, New York)	live	[150] [15] [825] [145] [47]
Pixalate ad fraud protection and analytics	Private; investors named in Sept 2021: Western Technology Investment and Javelin Venture Partners ($22.7 million raised...	Pixalate, Inc.	Fraud: invalid-traffic and bot-security companies	Post-bid fraud, brand-safety and display-viewability analytics and pre-bid IVT blocking across CTV, mobile apps and websites, for agencies, DSPs, SSPs and publishers (buying task: IVT measurement and...	Scored in cohort B	yes (VER-U-S05: letter addressed to Santa Monica, California)	live	[24] [474] [151] [65] [743]
Fraudlogix IVT detection	Unconfirmed (founded 2010 and still led by its founders; no ownership statement)	Fraudlogix	Fraud: invalid-traffic and bot-security companies	Post-bid pixel analytics of served impressions (GIVT and SIVT, brand-safety flags, domain verification, viewability) and a flat-fee pre-bid IP blocklist, sold exclusively to high-volume RTB platforms...	Scored in cohort B	yes (VER-U-S19: Hallandale Beach, FL)	live	[33] [746] [745] [47] [65]
Escalated IVT scoring	Unconfirmed (founder and CEO Kyle Smith; no ownership statement)	Escalated Solutions, LLC	Fraud: invalid-traffic and bot-security companies	Pre-bid and post-bid GIVT and SIVT detection and blocking for display, video, app and CTV, sold across the sell side and buy side (buying task: IVT filtration)	Scored in cohort B	yes (VER-U-S23: Carmel, Indiana dateline)	live	[34] [652] [653] [65]
FouAnalytics	Unconfirmed (built and run by its founder per his FAQ)	FouAnalytics (Dr. Augustine Fou); legal entity not stated	Fraud: invalid-traffic and bot-security companies	JavaScript-tag measurement of bots and humans in ads and on sites, sold to large advertisers by annual subscription and free below stated volumes (buying task: IVT measurement for buyers)	Scored in cohort B	unconfirmed (no location on the pages opened)	live (sign-up by invitation; author's articles continue to Nov 2026)	[826] [36] [47]
Zefr brand safety and suitability	Unconfirmed (co-founded 2008 by Rich Raddon and Zach James; no ownership statement)	Zefr, Inc. (Los Angeles)	Contextual: content classification companies	Pre-bid content exclusions and post-campaign content-level reporting on YouTube, Meta, TikTok, Snap and other social platforms, sold as data solutions or managed media (buying task: suitability...	Scored in cohort C	yes (VER-U-S08: letter addressed to Los Angeles)	live	[37] [613] [827] [612] [493] [634]
Peer39 brand safety and suitability	Self-described 'Wholly independent, employee owned' since 2019; no filing	P39 Tech, LLC (Delaware)	Contextual: content classification companies	Pre-bid brand safety and suitability categories and tiers (page-level on the web, program-level on CTV) with post-buy reporting at no separate fee, for buyers on CTV, display and social (buying task...	Scored in cohort C	yes (VER-U-S34: Delaware LLC; VER-U-S35: integrations with US DSPs)	live	[16] [760] [521] [25] [47]

PRODUCT	OWNER	OPERATING ENTITY	PEER GROUP	WHAT IT SELLS	IN THIS PAPER	US AVAILABILITY	STATUS AT CUTOFF	SRC
Proximic by Comscore	Comscore, Inc	Comscore, Inc.	Contextual: content classification companies	Pre-bid brand safety, suitability, misinformation and media-bias filters across desktop, mobile, CTV and podcasts, for media buyers (buying task: suitability avoidance)	Scored in cohort C	yes (Comscore operates US services per VER-U-S58; the product page VER-U-S57...	live	[17] [633] [47]
illumaSentinel	Unconfirmed	Illuma Technology Limited (England and Wales 06727425)	Contextual: content classification companies	Pre-bid page-level brand safety and suitability classification with reporting, activated in The Trade Desk's brand-safety marketplace (buying task: suitability avoidance)	Scored in cohort C	yes (VER-U-S61: operates from London and New York; VER-U-S60: available in The...	live	[125] [39] [753] [47]
Mantis Sentinel	Reach plc (FAQ: Reach developed Mantis in 2019)	Mantis (Reach plc)	Contextual: content classification companies	Advertiser tag that records every URL an ad is served on, aggregates impressions at page level and reports brand-safety insight (buying task: post-serve brand-safety reporting)	Scored in cohort C	unconfirmed	live (offered on the advertiser page; no launch date)	[751] [750] [47]
Mantis Contextual and Mantis Video	Reach plc	Mantis (Reach plc)	Contextual: content classification companies	Publisher-side page and video brand-safety and contextual classification, charged per URL scanned and passed to ad servers as key-values for PMP and PG deals (buying task: a seller classifying its...	Listed, not scored	unconfirmed	live (Mantis Video launched 2 Jun 2026)	[750] [752] [751]
ViewIQ	Private: Truelink Capital completed a significant investment on 24 Apr 2025 (terms not disclosed); Star Mountain...	Channel Factory (City of Industry, CA)	Contextual: content classification companies	YouTube pre-bid content classification with inclusion and exclusion lists, sold with Channel Factory's platform-native media buying on YouTube and social (buying task: managed media with suitability)	Scored in cohort C	yes (VER-U-S54: corporate headquarters, City of Industry, CA)	live	[19] [40] [540] [47]
Pixability brand suitability	Unconfirmed	Pixability	Contextual: content classification companies	YouTube brand-suitability controls, contextual segments and AI curation inside Pixability's YouTube outcomes platform and media services (buying task: YouTube media and optimisation)	Scored in cohort C	unconfirmed (no address on the pages opened)	live	[819] [818] [47]
Sightly Brand Mentality	Unconfirmed	Sightly	Contextual: content classification companies	YouTube and programmatic media buying with AI-based brand-suitability controls (buying task: media with embedded suitability)	Scored in cohort C	unconfirmed	live	[823] [47]
Apostra Brand Standards	Unconfirmed	Apostra, Inc. (formerly Scope3)	Contextual: content classification companies	LLM-based buy or don't-buy brand-suitability decisions for agent-driven buying, with a measurement-only option, as documented by Scope3 in 2025	Listed, not scored	unconfirmed	unconfirmed (documented in 2025; no current page or document on 10 Nov 2026; no retirement notice found)	[790] [828] [770] [321] [345] [47]
NewsGuard for Advertising	Unconfirmed	NewsGuard	Contextual: content classification companies	Source-level journalistic reliability ratings of news sites used for inclusion and avoidance in brand campaigns (buying task: news inclusion lists)	Listed, not scored	yes (VER-U-S95: New York address on the page)	live	[829] [47]

PRODUCT	OWNER	OPERATING ENTITY	PEER GROUP	WHAT IT SELLS	IN THIS PAPER	US AVAILABILITY	STATUS AT CUTOFF	SRC
IRIS.TV video data connectivity	Unconfirmed	IRIS.TV	Contextual: content classification companies	Video-level metadata connectivity for contextual targeting and partners' products (buying task: data pipe)	Listed, not scored	unconfirmed	live	[830] [47]
GumGum contextual advertising	Unconfirmed	GumGum, Inc.	Contextual: content classification companies	Contextual media with pre-bid suitability exclusion built into GumGum's own media (buying task: contextual media)	Listed, not scored	unconfirmed	live (the Verity name is no longer used; gumgum.com/verity redirects)	[43] [47]
Seedtag	Unconfirmed	Seedtag	Contextual: content classification companies	Contextual media; brand-safety filters and publisher block and allow lists applied inside its own campaigns	Listed, not scored	unconfirmed	live	[44] [47]
Vobile content protection	Vobile Group Limited (listing not re-checked here)	Vobile Group Limited	Other: ad quality and content protection	Content protection and rights management for rights holders	Listed, not scored	unconfirmed	live	[831] [47]
TrafficGuard	Adveritas Limited (ASX: AV1)	TrafficGuard	Fraud: invalid-traffic and bot-security companies	Click, install and conversion fraud protection for Google Search, Performance Max, Meta, affiliate and mobile app campaigns, with impression validation only as plan features and an advertiser Ad Tag...	Listed, not scored	yes (VER-U-S29: US operations formally launched, 18 Feb 2026)	live	[832] [833] [834] [835] [836] [837]
Anura Ad Tag Protect	Unconfirmed	Anura Solutions LLC	Fraud: invalid-traffic and bot-security companies	A publisher's AdSense ad-tag wrapper that withholds ads from flagged visitors; visitor-level fraud scoring on advertisers' landing pages	Listed, not scored	yes (VER-U-S74: Middletown, Delaware address)	live	[41] [42] [47]
CHEQ	Unconfirmed	CHEQ AI Technologies Ltd.	Fraud: invalid-traffic and bot-security companies	Click, lead, bot and audience protection for paid media and websites	Listed, not scored	yes (VER-U-S76: New York office)	live	[838] [47]
Lunio	Unconfirmed	Lunio	Fraud: invalid-traffic and bot-security companies	Click-based invalid-traffic detection with placement and audience exclusions for search, social, display and Performance Max	Listed, not scored	unconfirmed (the 9 Nov check's Chicago office was not re-opened)	live	[839] [47]
Spider AF	Unconfirmed	Spider AF	Fraud: invalid-traffic and bot-security companies	Click and invalid-traffic blocking for web advertising, plus lead and affiliate fraud tools	Listed, not scored	unconfirmed	live	[840] [47]
fraud0	Unconfirmed	fraud0	Fraud: invalid-traffic and bot-security companies	Invalid-click and billing-error detection with automated refund claims to Google Ads	Listed, not scored	unconfirmed	live	[841] [47]
AppsFlyer Fraud Protection	Unconfirmed	AppsFlyer	Fraud: invalid-traffic and bot-security companies	Mobile install and in-app event fraud protection for app marketers	Listed, not scored	unconfirmed	live	[842] [47]
GeoEdge	Unconfirmed	GeoEdge	Other: ad quality and content protection	Ad quality and malvertising protection for publishers and platforms	Listed, not scored	unconfirmed	live	[843] [47]
Confiant	Unconfirmed	Confiant	Other: ad quality and content protection	Blocking of malicious and low-quality ads for publishers and platforms	Listed, not scored	unconfirmed	live	[844] [47]

PRODUCT	OWNER	OPERATING ENTITY	PEER GROUP	WHAT IT SELLS	IN THIS PAPER	US AVAILABILITY	STATUS AT CUTOFF	SRC
The Media Trust	Unconfirmed	The Media Trust	Other: ad quality and content protection	Real-time blocking of malware, malvertising and offensive ads for publishers, ad-tech platforms and app developers	Listed, not scored	unconfirmed	live	[845] [47]
xpln.ai attention measurement	Unconfirmed	xpln.ai	Measurement: audience and media measurement companies	Attention metrics from eye-tracking panel data and machine learning for advertisers, agencies and marketplaces	Listed, not scored	unconfirmed	live	[846] [65]
Oracle Moat	Oracle Corporation	Oracle	Verification: third-party verification companies	Viewability, invalid-traffic and attention measurement (retired)	Historical	no (retired)	retired (30 Sep 2024)	[45] [47]
Oracle Contextual Intelligence	Oracle Corporation	Oracle	Contextual: content classification companies	Pre-bid contextual and brand-safety segments (retired)	Historical	no (retired)	retired (30 Sep 2024)	[45] [47]
OpenSlate	DoubleVerify Holdings, Inc	OpenSlate (acquired by DoubleVerify)	Contextual: content classification companies	YouTube brand suitability and content classification (brand retired)	Historical	no (brand retired)	retired brand (absorbed)	[754] [755] [757]
Meetrics	DoubleVerify Holdings, Inc	Meetrics GmbH (acquired by DoubleVerify)	Verification: third-party verification companies	Viewability and ad verification in Europe (absorbed)	Historical	no (absorbed)	retired brand (absorbed)	[754] [757]
Adverif.ai	Zefr, Inc	Adverif.ai (acquired by Zefr)	Contextual: content classification companies	AI misinformation identification (absorbed into Zefr)	Historical	no (absorbed)	retired brand (absorbed)	[488]
Forensiq	Impact (impact.com) (acquired 14 Jun 2016)	Forensiq (acquired by Impact Radius)	Fraud: invalid-traffic and bot-security companies	Advertising fraud prevention (absorbed into Impact)	Historical	no (absorbed)	retired brand (absorbed)	[756]
Kubient KAI	Kubient, Inc	Kubient, Inc.	Fraud: invalid-traffic and bot-security companies	Pre-bid ad-fraud detection deployable on third-party RTB platforms, for advertisers and publishers (historical)	Historical	unconfirmed	unconfirmed (company delisted; no current offer found)	[747] [748] [749]
Method Media Intelligence	Unconfirmed	Method Media Intelligence	Fraud: invalid-traffic and bot-security companies	Unsampled impression measurement for IVT, viewability and brand safety across display, video, CTV, audio and DOOH (historical)	Historical	unconfirmed	unconfirmed (site repurposed; no closure notice found)	[46]
Comscore vCE	Comscore, Inc.	Comscore, Inc.	Verification: third-party verification companies	Campaign validation and audience measurement (historical)	Historical	unconfirmed	unconfirmed (no current product page; no retirement notice)	[633]

Downloads

The datasets expose the source and claim registers, unit grid, scoring rules, cells and derived checks. Read the status dates and access limits alongside the data. The December freeze is a scheduled recheck, not a completed review in this November draft.

Downloads

The data files are in the data/ folder beside this PDF, in the package and on the web page. Open the HTML edition to use these as links.

The two other editions, then the data as CSV and JSON, UTF-8. The data files are the canonical files every chart, matrix and count in the paper was built from.

1. ad-verification-2026.pdf — the print edition: A4, with page numbers and chapter bookmarks
2. ad-verification-2026.md — the reading edition in Markdown, with numbered references
3. data/source-register.csv — every source with publication, event and access dates, type, geography, funding, access limits and reliability notes
4. data/claim-ledger.csv — every material claim with sources, passage locator, scope, evidence class, review outcome, confidence and final wording
5. data/capability-matrix-wide.csv — vendors × dimensions, scores only (0–5, n/e, n/a)
6. data/capability-matrix-long.csv — one row per cell with availability, evidence grade, date, rationale and sources
7. data/capability-dimensions.csv — the dimension definitions and anchors
8. data/vq.csv — Verdict Quotient per vendor: group means, coverage, band, the reweightings and every dimension score
9. data/vq-method.csv — VQ groups, weights, shrinkage, field mean, band thresholds, and the reweighting, shrinkage and missing-evidence tests
10. data/vq-missing-evidence.csv — every scored cell turned into n/e one at a time, with the VQ before and after
11. data/swot.csv — every SWOT line for the profiled vendors with its basis (record, vendor documentation, vendor claim, inference) and sources
12. data/deep-profiles.csv — the deep-profile summaries: what is sold, buyers, earnings, dependencies, fit, limits
13. data/deep-profile-records.csv — every profiled capability, performance claim, financial figure and timeline event with sources
14. data/h1a-units.csv — the H1a unit grid: one row per product, metric family and environment sold, with its MRC status and sources
15. data/study-evidence.csv — studies with design, sample, sponsor and limits
16. data/corrections-ledger.csv — every change made after verification and challenge passes, with implementation and verification
17. data/vendor-universe.csv — the reasoned vendor sample with scope group, owner, formats, routes and status
18. data/agentic-map.csv — AI and agentic capabilities with status labels and observability
19. data/rules.csv — rules and governance instruments by jurisdiction

20	data/estimates.csv — market-size estimates with inclusions, exclusions and combinability
21	data/decision-tool.csv — the decision tool rows
22	data/product-register.csv — the product register: every scored or listed product with its purchased service, operating entity, owner, cohort, inclusion reason, aliases and US availability
23	data/hypotheses.csv — the hypotheses registered before extraction, with the rules for supported, contradicted and unresolved
24	data/figure-data.json — the canonical dataset behind every data figure
25	data/derived-checks.json — every derived number with its formula and inputs
26	data/stats.json — every count the prose and figures quote about the research itself, computed from the ledgers
27	data/scoring-rules.json — dimension definitions, anchors, cell rules and the VQ definition, fixed before scoring

REFERENCES

Numbered sources

The 1069 sources cited in this paper, numbered in order of first citation. The full register of 1458 source records, with access dates, access limits and reliability notes, is in data/source-register.csv. Click any number in the text to open the same record.

[1] Invalid Traffic Detection and Filtration Standards Addendum (IVT 2.0), June 2020 Update (Final). Media Rating Council (MRC), 2020-06. https://mediaratingcouncil.org/sites/default/files/Standards/IVT%20Addendum%20Update%20062520.pdf Accessed 2026-11-10.

[2] MRC Viewable Ad Impression Measurement Guidelines (Desktop). Media Rating Council (MRC), 2015-08-18. https://mediaratingcouncil.org/sites/default/files/Standards/081815%20Viewable%20Ad%20Impression%20Guideline_v2.0_Final.pdf Accessed 2026-11-10.

[3] MRC Supplement to IAB Guidelines for the Conduct of Ad Verification: Enhanced Content Level Context and Brand Safety. Media Rating Council (MRC) (sponsoring associations: 4A's, ANA, IAB), 2018-09. https://mediaratingcouncil.org/sites/default/files/Standards/MRC%20Ad%20Verification%20Supplement-%20Enhanced%20Content%20Level%20Context%20and%20Brand%20Safety%20%28Final%29.pdf Accessed 2026-11-10.

[4] AUTHOR'S WORK · Where the Surplus Settles. Evgeny Popov, No Fluff Advisory, 2026-09-25. https://nofluffadvisory.com/writing/where-the-surplus-settles/ Accessed 2026-11-11.

[5] AUTHOR'S WORK · The Risk You Can Price. Evgeny Popov, No Fluff Advisory, 2026-08-23. https://nofluffadvisory.com/writing/the-risk-you-can-price/ Accessed 2026-11-11.

[6] AgenticAdvertising.org registry API: list agents (no filter, and type=measurement, governance, signals, sales, creative, buying). AgenticAdvertising.org, 2026-11-10. https://agenticadvertising.org/api/registry/agents?type=measurement Accessed 2026-11-10.

[7] DV Media Quality: Verify and Protect. DoubleVerify, undated (©2026 DoubleVerify). https://doubleverify.com/products/advertisers/verify Accessed 2026-11-10.

[8] Ad Verification (solution page). Integral Ad Science, undated. https://integralads.com/solutions/ad-verification/ Accessed 2026-11-10.

[9] MRC Mobile Viewable Ad Impression Measurement Guidelines. Media Rating Council (MRC), 2016-06-28. https://mediaratingcouncil.org/sites/default/files/Standards/062816%20Mobile%20Viewable%20Guidelines%20Final.pdf Accessed 2026-11-10.

[10] Responsible Media Guide V1, Section 05: Brand Safety. ISBA (with The Aperto Partnership), 2025-03. https://www.isba.org.uk/media/responsible-media-guide-brand-safety Accessed 2026-11-10.

[11] Policy for Property-Level Ad Verification Representations. Media Rating Council (MRC), 2025-10-18. https://mediaratingcouncil.org/sites/default/files/News/Ad%20Verification%20Representation%20Policy%20102025.pdf Accessed 2026-11-10.

[12] Pixalate documentation: Pixalate Analytics (overview). Pixalate, Inc., undated (live page). https://docs.pixalate.com/analytics/overview Accessed 2026-11-10.

[13] IAB Tech Lab Expands Open Measurement SDK to New CTV Platforms (press release). IAB Tech Lab, 2024-05-30. https://iabtechlab.com/press-releases/iab-tech-lab-expands-open-measurement-sdk-to-new-ctv-platforms/ Accessed 2026-11-10.

[14] FraudSensor Overview (Advertising Protection documentation). HUMAN Security, 2026-06-29. https://docs.humansecurity.com/advertising/fraudsensor-overview Accessed 2026-11-10.

[15] MediaGuard Overview. HUMAN Security (documentation site), undated. https://docs.humansecurity.com/advertising/mediaguard-overview Accessed 2026-11-10.

[16] Third-Party Ad Verification: Pre-Bid Controls. Peer39, 2026. https://www.peer39.com/verification Accessed 2026-11-10.

[17] For Media Buyers - Proximic by Comscore. Comscore, Inc. (Proximic by Comscore), 2026. https://www.comscore.com/home/Programmatic-Targeting-Solutions/For-Media-Buyers Accessed 2026-11-10.

[18] Adloox for advertisers. Adloox, undated (footer © 2025). https://www.adloox.com/solutions/advertisers Accessed 2026-11-10.

[19] ViewIQ & ActivateIQ: YouTube Ad Intelligence (product page). Channel Factory, 2026. https://channelfactory.com/product Accessed 2026-11-10.

[20] Self-Service Media. Sightly Enterprises, Inc., undated (page footer: (c) 2026 Sightly Enterprises, Inc.). https://www.sightly.com/self-service-media/ Accessed 2026-11-10.

[21] Brand Safety Certified Guidelines. TAG Brand Safety Certification Working Group, 2026-07. https://www.tagtoday.net/hubfs/BSC%20Guidelines.pdf Accessed 2026-11-10.

[22] YouTube Self Service (Self-Service Activation, Supported by Pixability). Pixability, Inc., undated (sitemap lastmod 2026-11-03). https://www.pixability.com/youtube-solutions/advertising/self-serve Accessed 2026-11-10.

[23] Terms and Conditions. Channel Factory LLC, 2026-08-17. https://channelfactory.com/terms-conditions Accessed 2026-11-10.

[24] Pixalate home page (Ad Fraud Protection, Privacy, and Compliance Platform). Pixalate. https://www.pixalate.com/ Accessed 2026-11-10.

[25] Peer39 Acquires Adloox, Extending Brand Safety, Contextual Intelligence and Verification Across DV360, Meta, and Amazon DSP. Peer39, 2026-06-03 (event 2026-06-02). https://www.peer39.com/blog/peer39-acquires-adloox-extending-brand-safety-contextual-intelligence-and-verification-across-dv360-meta-and-amazon-dsp Accessed 2026-11-10.

[26] DoubleVerify Holdings, Inc. Form 10-K for the fiscal year ended 31 December 2025. DoubleVerify Holdings, Inc. (SEC filing), 2026-02-26. https://www.sec.gov/Archives/edgar/data/1819928/000110465926020499/dv-20251231×10k.htm Accessed 2026-11-09.

[27] Integral Ad Science Holding Corp. Form 10-K for the fiscal year ended 31 December 2024. Integral Ad Science Holding Corp. (SEC filing), 2025-02-28 (event 2024-12-31). https://www.sec.gov/Archives/edgar/data/1842718/000184271825000023/ias-20241231.htm Accessed 2026-11-09.

[28] Integral Ad Science Holding Corp. Form 8-K (closing of the merger). Integral Ad Science Holding Corp. (SEC filing), 2025-12-23. https://www.sec.gov/Archives/edgar/data/1842718/000114036125046475/ef20061707_8k.htm Accessed 2026-11-09.

[29] Video & CTV | Adloox. Adloox, undated (footer © 2025). https://www.adloox.com/channels/video Accessed 2026-11-10.

[30] Protected by Mediaocean (home page). Protected Media (Mediaocean), 2026. https://www.protected.media/ Accessed 2026-11-10.

[31] IVT Detection | Protected by Mediaocean. Protected by Mediaocean, undated (footer 2026). https://www.protected.media/ivt-detection Accessed 2026-11-10.

[32] Ad Traffic Validation For Branding Campaigns. mFilterIt, undated (footer: 2026 copyright). https://www.mfilterit.com/ad-traffic-validation-for-branding-campaigns Accessed 2026-11-10.

[33] IVT Detection: Pre-Bid & Post-Bid GIVT & SIVT Solutions (Programmatic IVT Detection product page). Fraudlogix, 2025-12-01 (event page metadata: published 2025-12-01, modified 2026-11-05). https://www.fraudlogix.com/ivt/ Accessed 2026-11-10.

[34] About Us (company page). Escalated Solutions, LLC, undated. https://escalated.io/company.html Accessed 2026-11-10.

[35] Features (plans and feature table; /pricing.html redirects here). Escalated Solutions, LLC, undated. https://escalated.io/features.html Accessed 2026-11-10.

[36] Frequently Asked Questions about FouAnalytics. Dr. Augustine Fou (founder), 2022-12-05. https://www.linkedin.com/pulse/frequently-asked-questions-fouanalytics-augustine-fou Accessed 2026-11-10.

[37] Brand Safety & Suitability (Zefr solution page). Zefr, 2026-03-25. https://zefr.com/solution/brand-safety-suitability Accessed 2026-11-10.

[38] Connected TV (CTV) Contextual Targeting. Peer39, undated (footer © 2026). https://www.peer39.com/connected-tv Accessed 2026-11-10.

[39] Unlock brand safe inventory at scale with illuma + The Trade Desk. Illuma Technology Limited, 2026-07-27 (event 2026-07-27 (availability in The Trade Desk's brand safety marketplace announced)). https://weareilluma.com/brand-safety-on-the-trade-desk/ Accessed 2026-11-10.

[40] Truelink Capital Completes Investment in Channel Factory to Fuel Global Expansion. Channel Factory and Truelink Capital (GlobeNewswire), 2025-04-24. https://www.globenewswire.com/news-release/2025/04/24/3067595/0/en/truelink-capital-completes-investment-in-channel-factory-to-fuel-global-expansion.html Accessed 2026-11-09.

[41] Protect Your Ad Tags from Fraud (Ad Tag Protect). Anura Solutions LLC, undated (© 2026). https://www.anura.io/ad-tag-protect Accessed 2026-11-10.

[42] Ad Fraud Detection & Prevention Software (FAQ). Anura Solutions LLC, undated. https://www.anura.io/ad-fraud-detection Accessed 2026-11-10.

[43] Contextual Advertising (target of https://gumgum.com/verity). GumGum, Inc., undated. https://gumgum.com/contextual-advertising Accessed 2026-11-10.

[44] Brand Safety Policy. Seedtag, undated. https://www.seedtag.com/brand-safety-policy Accessed 2026-11-10.

[45] Oracle Advertising EOL Frequently Asked Questions. Oracle Corporation, 2024-08-12 (event 2024-09-30). https://www.oracle.com/a/ocom/docs/oracle-advertising-end-of-life-faqs.pdf Accessed 2026-11-09.

[46] methodmi.com (live page). Unknown operator (page titled Method Media Intelligence), undated. https://methodmi.com/ Accessed 2026-11-10.

[47] Accredited Services: Digital (metrics by service). Media Rating Council (MRC), 2026. https://mediaratingcouncil.org/accreditation/digital Accessed 2026-11-10.

[48] 2024 IVT Interim Updates (memo to digital measurement vendors). Media Rating Council (MRC), 2024-04-24. https://mediaratingcouncil.org/sites/default/files/Standards/2024_IVT_Interim_Updates_FINAL.pdf Accessed 2026-11-10.

[49] Server-Side Ad Insertion and OTT Guidance. Media Rating Council (MRC), 2021-08. https://mediaratingcouncil.org/sites/default/files/Standards/083021%20SSAI%20and%20OTT%20Guidance%20%20FINAL.pdf Accessed 2026-11-10.

[50] Accreditation letter to Protected Media: property-level ad verification (GI092525). Media Rating Council (MRC), 2025-09-25. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Protected%20Media%20AV/GI092525_Protected%20Media%20Ad%20Verification%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[51] TAG Registry search results (registry widget at tagtoday.net/registry, served from hubgloo.opxar.com). TAG, 2026. https://hubgloo.opxar.com/search?q=DoubleVerify Accessed 2026-11-10.

[52] Standards & Guidelines (list of MRC measurement standards). Media Rating Council (MRC), 2026. https://mediaratingcouncil.org/standards-and-guidelines Accessed 2026-11-10.

[53] MRC accreditation decisions report: Q1-23 (January-March 2023). Media Rating Council (MRC), 2023-04-25. https://mediaratingcouncil.org/sites/default/files/News/Accreditation-Matters/042523_MRC_Accreditation_Status_Communication.pdf Accessed 2026-11-10.

[54] MRC accreditation decisions report: Q4-23 (October-December 2023). Media Rating Council (MRC), 2024-01-08. https://mediaratingcouncil.org/sites/default/files/News/Accreditation-Matters/010824_MRC_Accreditation_Status_Communication.pdf Accessed 2026-11-10.

[55] MRC accreditation decisions report: Q2-24 (April-June 2024). Media Rating Council (MRC), 2024-07-26. https://mediaratingcouncil.org/sites/default/files/News/MRC%20Accreditation%20Status%20Quarterly%20Update_Q2%2024%207.26.pdf Accessed 2026-11-10.

[56] MRC accreditation decisions report: Q3-25 (July-September 2025). Media Rating Council (MRC), 2025-10-16. https://mediaratingcouncil.org/sites/default/files/News/MRC%20Accreditation%20Status%20Quarterly%20Update_Q3%2025_10.15.25%20.pdf Accessed 2026-11-10.

[57] MRC accreditation decisions report: Q2-23 (April-June 2023). Media Rating Council (MRC), 2023-07-14. https://mediaratingcouncil.org/sites/default/files/News/Accreditation-Matters/071423_MRC_Accreditation_Status_Communication.pdf Accessed 2026-11-10.

[58] MRC accreditation decisions report: Q3-23 (July-September 2023). Media Rating Council (MRC), 2023-10-16. https://mediaratingcouncil.org/sites/default/files/News/Accreditation-Matters/101623_MRC_Accreditation_Status_Communication.pdf Accessed 2026-11-10.

[59] MRC accreditation decisions report: Q1-24 (January-March 2024). Media Rating Council (MRC), 2024-04-24. https://mediaratingcouncil.org/sites/default/files/News/050624_MRC_Accreditation_Status_Communication_FINAL.pdf Accessed 2026-11-10.

[60] MRC accreditation decisions report: Q2-25 (April-June 2025). Media Rating Council (MRC), 2025-07-11. https://mediaratingcouncil.org/sites/default/files/News/MRC%20Accreditation%20Status%20Quarterly%20Update_Q2%2025_7.9.25.pdf Accessed 2026-11-10.

[61] MRC accreditation decisions report: Q4-25 (October-December 2025). Media Rating Council (MRC), 2026-01-22. https://mediaratingcouncil.org/sites/default/files/News/MRC%20Accreditation%20Status%20Quarterly%20Update%20Q4%202025.pdf Accessed 2026-11-10.

[62] MRC accreditation decisions report: Q1-26 (January-March 2026). Media Rating Council (MRC), 2026-04-17. https://mediaratingcouncil.org/sites/default/files/News/MRC%20Accreditation%20Status%20Quarterly%20Update_Q1%202026.pdf Accessed 2026-11-10.

[63] MRC accreditation decisions report: Q2-26 (April-June 2026). Media Rating Council (MRC), 2026-07-13. https://mediaratingcouncil.org/sites/default/files/News/MRC%20Accreditation%20Status%20Quarterly%20Update_Q2%2026%207.14.26.pdf Accessed 2026-11-10.

[64] MRC accreditation decisions report: Q3-26 (July-September 2026). Media Rating Council (MRC), 2026-10-16. https://mediaratingcouncil.org/sites/default/files/News/Accreditation-Matters/MRC_Accreditation_Status_Quarterly_Update_Q3_2026.pdf Accessed 2026-11-10.

[65] Services Under Review: Digital. Media Rating Council (MRC), 2026. https://www.mediaratingcouncil.org/accreditation/services-under-review/digital Accessed 2026-11-10.

[66] Accreditation letter to Integral Ad Science: S2S integration with Amazon DSP (GI110725). Media Rating Council (MRC), 2025-11-07. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/IAS/GI110725_IAS%203P%20Amazon%20DSP%20Accreditation%20Letter%20.pdf Accessed 2026-11-10.

[67] Accreditation letter to Integral Ad Science (GI070926: Campaign Monitor, Firewall and Viewability Systems; Meta and Google ADH integrations). Media Rating Council (MRC), 2026-07-09 (event 2026-06 (Board decision per S31)). https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/IAS/GI070926_Integral%20Ad%20Science%20Recurring_Google%20_Meta%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[68] Accreditation letter to Google: YouTube content-level brand safety and suitability (GI052826). Media Rating Council (MRC), 2026-05-28. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Google/GI052826%20Google%20YouTube%20Content%20Level%20Brand%20Safety%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[69] ANA Programmatic Media Supply Chain Transparency Study: Complete Report (Playbook Summary pages 20-26). Association of National Advertisers (ANA); project team Reed Smith, Fiducia, TAG TrustNet, Lemonade Projects, Kroll, 2023-12 (event 2023-12-05 (ANA release of the report)). https://web-cdn.adslot.com/wp-content/uploads/2023/12/ana-programmatic-media-supply-chain-transparency-study.pdf Accessed 2026-11-10.

[70] Programmatic Transparency Benchmark: Q3 2025 Benchmark Findings. ANA, TAG TrustNet, Fiducia (brand-risk data by Mobian), 2025-11 (event 2025-11-05 (ANA release, per the brief-check notes)). https://www.fiducia.eco/_files/ugd/de006c_6a56fe29acbf4a609b81d0af2b7c5fc5.pdf Accessed 2026-11-10.

[71] Programmatic Transparency Benchmark: Q1 2026 Findings. ANA, TAG TrustNet, Fiducia, 2026-05 (event 2026-05-27 (ANA release)). https://www.fiducia.eco/_files/ugd/de006c_21911c212fab454bbd16b20517f333ae.pdf Accessed 2026-11-10.

[72] Programmatic Transparency Benchmark: Q4 2025 Benchmark Findings. ANA, TAG TrustNet, Fiducia, 2026. https://www.fiducia.eco/_files/ugd/de006c_730e6f3f60554d2983e0b75321eac55d.pdf Accessed 2026-11-10.

[73] ANA Programmatic Transparency Benchmark: Q2 2026 Working-Media Efficiency Reaches a Record 45% (repost of the ANA release). Fiducia (reposting the ANA), 2026-08-06. https://www.fiducia.eco/post/ana-programmatic-transparency-benchmark-q2-2026-working-media-efficiency-reaches-a-record-45 Accessed 2026-11-10.

[74] ANA/TAG/Fiducia Analysis Quantifies Level of 'AI Slop' in Digital Advertising Supply Chain for First Time (release, reposted by Fiducia). ANA, TAG (TAG TrustNet) and Fiducia; analysis drafted by Scott Cunningham (Cunningham.tech Consulting); reposted by Fiducia, 2026-08-10 (event release on or about 28 July 2026 (Adweek's article of that date says 'published Tuesday'; see VER-V-B-S02)). https://www.fiducia.eco/post/ana-tag-fiducia-analysis-quantifies-level-of-ai-slop-in-digital-advertising-supply-chain-for-first Accessed 2026-11-10.

[75] Are YouTube Advertisers Inadvertently Harvesting Data From Millions of Children?. Adalytics Research LLC, 2023-08-17. https://adalytics.io/blog/are-youtube-ads-coppa-compliant Accessed 2026-11-10.

[76] Does a lack of transparency create brand safety concerns for search advertisers?. Adalytics Research LLC, 2023-11. https://adalytics.io/blog/search-partners-transparency Accessed 2026-11-10.

[77] Are user IDs declared consistently in ad auctions?. Adalytics Research LLC, 2024-05-10. https://adalytics.io/blog/user-id-rotation Accessed 2026-11-10.

[78] Publisher Optimization Implementation Guide (IAS Help Center). Integral Ad Science, 2024-07-11 (event archived 2024-11-04). https://web.archive.org/web/20241104234011id_/https://helpcenter.integralplatform.com/article/publisher-optimization-implementation-guide Accessed 2026-11-10.

[79] IAS's Unwavering Commitment to Fostering a Safe Digital Advertising Ecosystem. Integral Ad Science, 2025-02-14. https://integralads.com/insider/ias-unwavering-commitment-to-foster-a-safe-digital-advertising-ecosystem/ Accessed 2026-11-10.

[80] IAS's Multi-Tiered Approach to Combating Ad Fraud (IAS Insider). Integral Ad Science, 2025-03-21. https://integralads.com/insider/iass-multi-tiered-approach-to-combating-ad-fraud/ Accessed 2026-11-10.

[81] Statement on Adalytics Report (copy archived by the Internet Archive on 7 February 2025, 21:24 UTC). DoubleVerify (archived by Internet Archive), 2025-02-07. https://web.archive.org/web/20250207212403id_/https://doubleverify.com/statement-on-adalytics-report/ Accessed 2026-11-10.

[82] MRC IVT Requirements and Processes for Pre-Bid: Comments on the Recent Adalytics Blog. Media Rating Council (MRC), 2025-04-11. https://mediaratingcouncil.org/sites/default/files/News/MRC%20Statement%20on%20pre%20bid%20IVT%20requirements%20and%20processes.pdf Accessed 2026-11-10.

[83] Inconsistencies in Classification of Online News Articles: A Call for Common Standards in Brand Safety Services (arXiv:2601.01303v1). Michael E. Smith, Riley Grossman, Antonio Torres-Agüero, Pritam Sen, Cristian Borcea, Yi Chen (NJIT; Torres-Agüero DeepSee.io), 2026-01-03. https://arxiv.org/abs/2601.01303 Accessed 2026-11-10.

[84] Quantifying the Cost of Ad Fraud: 2023-2028 (research report whitepaper). Juniper Research, 2023-09. https://fraudblocker.com/wp-content/uploads/2023/09/Ad-Fraud-Whitepaper_Juniper-Research.pdf Accessed 2026-11-09.

[85] Study Finds Ad Industry Anti-Fraud Programmes Saved European Advertisers EUR 3.45 Billion in 2023 (2025 European Ad Fraud Savings Report release). TAG (Trustworthy Accountability Group), 2025-06-17. https://www.tagtoday.net/pressreleases/study-finds-ad-industry-anti-fraud-programmes-saved-european-advertisers-3.45-billion-in-2023 Accessed 2026-11-10.

[86] Programmatic Ad Exchange Quality. Index Exchange. https://www.indexexchange.com/platform/exchange-quality/ Accessed 2026-11-10.

[87] Marketplaces Terms of Use. Index Exchange, 2026-05-13. https://www.indexexchange.com/marketplaces-terms-of-use/ Accessed 2026-11-10.

[88] OpenPath Publisher Terms and Conditions v.10.31.24. The Trade Desk, Inc., 2024-11. https://assets.thetradedesk.com/assets/documents/OpenPath-Publisher-Terms-and-Conditions-v.10.31.24.pdf Accessed 2026-11-10.

[89] Empowering businesses for a future that is conversational, personal, and agentic. Kya Sainsbury-Carter, Microsoft Advertising (blog), 2025-05-14 (event 2026-02-28). https://about.ads.microsoft.com/en/blog/post/may-2025/empowering-businesses-for-a-future-that-is-conversational-personal-and-agentic Accessed 2026-11-10.

[90] Microsoft Invest - Release notes. Microsoft (Microsoft Learn), 2025-09-23. https://learn.microsoft.com/en-us/xandr/invest/release-notes Accessed 2026-11-10.

[91] Addendum to the 4A's/IAB Standard Terms and Conditions for Internet Advertising (Media Buys). OpenX. https://assets.openx.com/downloads/legal/OpenX-Addendum-to-Standard-IAB-Agreement-(Media-Buys).pdf Accessed 2026-11-10.

[92] ANA Programmatic Transparency Benchmark Reveals Widening Gap Between Top and Bottom Advertisers. ANA, 2026-05-27. https://www.ana.net/content/show/id/pr-2026-05-programmatic Accessed 2026-11-10.

[93] Search upcoming agreements. Government Commercial Agency (GCA). https://www.gca.gov.uk/agreements/upcoming Accessed 2026-11-10.

[94] TAG Registry search (tagtoday.net/registry-search; endpoint hubgloo.opxar.com/search) for 'FouAnalytics', 'Fou' and 'Marketing Science Consulting'. Trustworthy Accountability Group (TAG), live search. https://www.tagtoday.net/registry-search Accessed 2026-11-10.

[95] Sizmek DSP, Inc. v. Integral Ad Science, Inc., Adv. No. 20-01134 (Bankr. S.D.N.Y.), docket. Free Law Project (CourtListener; PACER/RECAP), 2020-06-03 (event terminated 2021-04-12). https://www.courtlistener.com/docket/17219121/sizmek-dsp-inc-v-integral-ad-science-inc/ Accessed 2026-11-10. *Access: partial extract.*

[96] Accreditation letter to LinkedIn: Marketing Solutions platform. Media Rating Council (George W. Ivie), 2026-05-28. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/LinkedIn/GI052826_LinkedIn%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[97] Industry Leadership and Partnerships (X Help Center, Brand Safety). X Corp.. https://help.x.com/en/business-and-advertising/brand-safety/industry-leadership-and-partnerships Accessed 2026-11-10.

[98] Products That Protect (X Help Center, Brand Safety). X Corp.. https://help.x.com/en/business-and-advertising/brand-safety/products-that-protect Accessed 2026-11-10.

[99] Snap, IAS Brand Safety & Suitability Measurement Solution Now Available. Snap Inc. (Snapchat for Business blog), 2024-06-17. https://forbusiness.snapchat.com/blog/snap-ias-solution-ga Accessed 2026-11-10.

[100] Snap Launches New Brand Suitability Suite for Advertisers. Snap Inc. (Snapchat for Business blog), 2025-02-24. https://forbusiness.snapchat.com/blog/brand-suitability-suite-2025 Accessed 2026-11-10.

[101] Accreditation letter to DoubleVerify (Pinnacle platform; GI041326). Media Rating Council (MRC), 2026-04-13 (event 2026-04 (Board decision per S31)). https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/DoubleVerify/GI041326_DoubleVerify%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[102] CTV Programmatic Guide. IAB Tech Lab, 2025-12-11 (event modified 2026-01-30). https://iabtechlab.com/standards/advanced-tv/ctv-programmatic-guide/ Accessed 2026-11-10.

[103] Expanding Transparency in Streaming TV: A Call to Action for Our Industry (blog). DoubleVerify, 2025-12-04. https://doubleverify.com/blog/ctv/verify/expanding-transparency-in-streaming-tv-a-call-to-action-for-our-industry Accessed 2026-11-10.

[104] IAB Tech Lab Launches Device Attestation Support in Open Measurement SDK to Combat Device Spoofing. IAB Tech Lab, 2025-10-23. https://iabtechlab.com/press-releases/device-attestation-support-in-open-measurement-sdk/ Accessed 2026-11-10.

[105] IAB Tech Lab Announces OM SDK Support for Emerging CTV Ad Formats. IAB Tech Lab, 2026-11-10. https://iabtechlab.com/press-releases/iab-tech-lab-announces-om-sdk-support-for-emerging-ctv-ad-formats/ Accessed 2026-11-10.

[106] OM SDK Compliance API, 'latest' endpoint (certified integration partners and their latest compliance data). IAB Tech Lab. https://complianceomsdkapi.iabtechlab.com/compliance/latest Accessed 2026-11-10.

[107] IAB Tech Lab Announces CTV Ad Portfolio and Updated Guide to Programmatic CTV. IAB Tech Lab, 2025-12-11 (event public comment to 2026-01-31). https://iabtechlab.com/press-releases/iab-tech-lab-announces-ctv-ad-portfolio/ Accessed 2026-11-10.

[108] HUMAN (Formerly White Ops) ... Announces Discovery and Disruption of 'PARETO' CTV Botnet. HUMAN, 2021-04-21 (event found in 2020). https://web.archive.org/web/20210421155246id_/https://www.humansecurity.com/newsroom/human-formerly-white-ops-together-with-newly-formed-human-collective-and-industry-leaders-google-roku-announces-discovery-and-disruption-of-pareto-ctv-botnet Accessed 2026-11-10.

[109] DoubleVerify Unmasks CTV Ad Fraud Scheme, CycloneBot (release PDF on DV's investor site). DoubleVerify, 2024-02-28. https://s206.q4cdn.com/961864615/files/doc_news/2024/02/2024-02-28_DoubleVerify_Unmasks_CTV_Ad_Fraud_Scheme_280.pdf Accessed 2026-11-10.

[110] IAB Tech Lab Authorized Sellers for Apps (app-ads.txt) specification. IAB Tech Lab (authors from Google and Sharethrough), 2019-03. https://raw.githubusercontent.com/InteractiveAdvertisingBureau/Supply-Chain-Validation/main/app-ads.txt.md Accessed 2026-11-10.

[111] Authorized Digital Sellers with ads.txt and app-ads.txt (About page and version history). IAB Tech Lab, 2026-09-14. https://iabtechlab.com/ads-txt/ Accessed 2026-11-10.

[112] VAST CTV Addendum 2024. IAB Tech Lab Advanced TV Technical Working Group (GitHub: InteractiveAdvertisingBureau/vast), 2024 (event GitHub file created 2026-02-25). https://raw.githubusercontent.com/InteractiveAdvertisingBureau/vast/master/VAST%20CTV%20Addendum%202024.md Accessed 2026-11-10.

[113] DoubleVerify Extends Leadership in Social Video & CTV with Agreement To Acquire OpenSlate (Exhibit 99.2 to DoubleVerify's 8-K of 9 November 2021). DoubleVerify Holdings, Inc. (release filed with the SEC), 2021-11-09. https://www.sec.gov/Archives/edgar/data/1819928/000155837021015297/dv-20211109xex99d2.htm Accessed 2026-11-10.

[114] Integral Ad Science Holding Corp. Definitive Information Statement (Schedule 14C, merger). Integral Ad Science Holding Corp. (SEC filing), 2025-11-07 (event 2025-09-24). https://www.sec.gov/Archives/edgar/data/1842718/000114036125040934/ny20056670x2_defm14c.htm Accessed 2026-11-09. *Access: partial extract.*

[115] DoubleVerify Holdings, Inc. Form 8-K (Item 1.01, Merger Agreement). DoubleVerify Holdings, Inc. (SEC filing), 2026-08-10 (event 2026-08-06). https://www.sec.gov/Archives/edgar/data/1819928/000110465926092934/tm2621904d5_8k.htm Accessed 2026-11-10.

[116] Flashtalking Acquires Protected Media, Takes Position as Only Independent Ad Server with MRC Accredited Solution for Fraud Detection in the OTT Segment. Flashtalking, 2021-02-23 (event 2021-02). https://web.archive.org/web/20211103203244id_/https://www.flashtalking.com/blog/2021/2/22/flashtalking-acquires-protected-media-takes-position-as-only-independent-ad-server-with-mrc-accreditation-for-fraud-detection-in-the-ott-segment Accessed 2026-11-09.

[117] 4A's Advertiser Protection Bureau Delivers Brand Suitability Framework and Brand Safety Floor. 4A's (via PR Newswire), 2018-09-21. https://www.prnewswire.com/news-releases/4as-advertiser-protection-bureau-delivers-brand-suitability-framework-and-brand-safety-floor-in-move-to-help-advertisers-assess-risk-300717002.html Accessed 2026-11-10.

[118] 4As APB Media Responsibility Principles. 4A's (Media, Tech and Data practice with the Advertiser Protection Bureau), 2023-12-11. https://www.aaaa.org/blog/4as-apb-media-responsibility-principles/ Accessed 2026-11-10.

[119] Brand Safety Floor + Suitability Framework (hosted copy). Brand Safety Institute (BSI). https://www.brandsafetyinstitute.com/resources/frameworks/brand-safety-floor-suitability Accessed 2026-11-10.

[120] Understanding Brand Safety & Brand Suitability in a Contemporary Media Landscape. IAB, 2020-12. https://www.iab.com/wp-content/uploads/2020/12/IAB_Brand_Safety_and_Suitability_Guide_2020-12.pdf Accessed 2026-11-10.

[121] Responsible advertising industry initiatives – repository (Ad Responsibility Roundtable mapping). European Advertising Standards Alliance, 2023-12-11. https://www.easa-alliance.org/wp-content/uploads/2023/12/ITEM-4-Ad-Responsibility-Roundtable-mapping-of-initiatives-EX-SEP-23.pdf Accessed 2026-11-10.

[122] SAFER framework for digital brand safety. UK Government Communication Service, 2026-03-02. https://www.communications.gov.uk/publication/safer-framework-for-digital-brand-safety/ Accessed 2026-11-10.

[123] Content Taxonomy 2.2 (TSV). IAB Tech Lab (GitHub), 2020-12. https://github.com/InteractiveAdvertisingBureau/Taxonomies/blob/develop/Content%20Taxonomies/Content%20Taxonomy%202.2.tsv Accessed 2026-11-10.

[124] Brand Safety Playbook (4A's with TAG, BSI and IAB Tech Lab), hosted copy. Brand Safety Institute (hosting); 4A's et al. (authors), 2020-07. https://www.brandsafetyinstitute.com/hubfs/resource-library/brand-safety-playbook.pdf Accessed 2026-11-10.

[125] illuma Sentinel: Brand Safety & Suitability Technology. illuma, 2026. https://weareilluma.com/illuma-sentinel/ Accessed 2026-11-10.

[126] Our Brand Safety and Suitability Overview (IAS Insider). Integral Ad Science, 2024-03-06. https://integralads.com/insider/brand-safety-suitability-overview/ Accessed 2026-11-10.

[127] Work with Partners (Enterprise Marketing Portal). Google. https://enterprisemarketingportal.google/work-with-partners Accessed 2026-11-10.

[128] Hedge Fund Sues In Del. For Share Appraisal Damages. Law360 (Jeff Montgomery), 2026-02-26. https://law360.com/articles/2446561 Accessed 2026-11-10. *Access: partial extract.*

[129] Thor v. DoubleVerify Inc., docket 0:24-cv-00310 (D. Minn.). CourtListener (Free Law Project), 2024-02-05 (event 2025-01-27). https://www.courtlistener.com/docket/68280039/thor-v-doubleverify-inc/ Accessed 2026-11-10. *Access: partial extract.*

[130] Blackburn, Blumenthal Probe Amazon & Google After New Report Reveals They Placed Ads on Website that Hosts Child Sexual Abuse Material (press release). Office of U.S. Senator Marsha Blackburn, 2025-02-07. https://www.blackburn.senate.gov/2025/2/issues/technology/blackburn-blumenthal-probe-amazon-google-after-new-report-reveals-they-placed-ads-on-website-that-hosts-child-sexual-abuse-material Accessed 2026-11-10.

[131] X Corp. v. WFA, Fifth Circuit No. 26-10394, Stipulation of Dismissal as to X and WFA (Doc. 119). X Corp. and World Federation of Advertisers, 2026-08-05. https://storage.courtlistener.com/recap/gov.uscourts.ca5.229782/gov.uscourts.ca5.229782.119.0.pdf Accessed 2026-11-10.

[132] Letter from Senator Mark R. Warner to FTC Chairman Andrew Ferguson. U.S. Senator Mark R. Warner, 2025-03-24. https://www.warner.senate.gov/public/_cache/files/9/b/9beb62bb-9be4-4003-862c-7f31e5dc64e7/28124EF34207E1976F98496DE79C2CE9F214F1AE419D59DDA56205D51C396709.250324.warner-ferguson-bot-fraud-letter.final.signed.pdf Accessed 2026-11-10.

[133] DoubleVerify Holdings, Inc. Definitive Proxy Statement (Schedule 14A, merger). DoubleVerify Holdings, Inc. (SEC filing), 2026-09-29 (event 2026-10-29). https://www.sec.gov/Archives/edgar/data/1819928/000110465926111962/tm2621904-12_defm14a.htm Accessed 2026-11-09. *Access: partial extract.*

[134] Electrical Workers Pension Fund, Local 103, I.B.E.W. v. DoubleVerify Holdings, Inc., class action complaint (ECF 1), 1:25-cv-04332 (S.D.N.Y.). Plaintiff pension fund; Labaton Keller Sucharow LLP, 2025-05-22. https://storage.courtlistener.com/recap/gov.uscourts.nysd.643090/gov.uscourts.nysd.643090.1.0.pdf Accessed 2026-11-10.

[135] CourtListener RECAP party-name searches for the scored vendors (results pages). CourtListener (Free Law Project). https://www.courtlistener.com/?type=r&party_name=Doubleverify&order_by=dateFiled%20desc Accessed 2026-11-10.

[136] Liste des avis et décisions: searches 'adloox' (1 result) and 'vérification publicitaire' (2 results). Autorité de la concurrence (France), 2026. https://www.autoritedelaconcurrence.fr/fr/liste-des-decisions-et-avis?search_a

pi_fulltext=adloox Accessed 2026-11-10.

[137] Agencies look to AI agents for edge, but must fend off brands looking to in-house with the same tech. Sam Bradley / Digiday, 2026-06-24 (event 2025-12 to 2026-06). https://digiday.com/media-buying/agencies-look-to-ai-agents-for-edge-but-must-fend-off-brands-looking-to-in-house-with-the-same-tech/ Accessed 2026-11-11.

[138] Butler/Till's first agentic media buying tests cut media and supply chain costs. Sam Bradley / Digiday, 2026-03-17 (event 2025-12 to 2026-06). https://digiday.com/media-buying/butler-tills-first-agentic-media-buying-tests-cut-media-and-supply-chain-costs/ Accessed 2026-11-11.

[139] OpenProposal issues #126-#129 (implementation comments filed by the author). Evgeny Popov (the author), GitHub IABTechLab/OpenProposal, 2026-10-27. https://github.com/IABTechLab/OpenProposal/issues/126 Accessed 2026-11-10.

[140] OpenProposal README (main branch), status table. IAB Tech Lab (IABTechLab/OpenProposal on GitHub), 2026-09-22. https://github.com/IABTechLab/OpenProposal/blob/main/README.md Accessed 2026-11-10.

[141] IABTechLab/seller-agent branch feat/openproposal-mapping (src/ad_seller/services/openproposal_service.py; src/ad_seller/models/openproposal.py). IAB Tech Lab (AAMP seller agent on GitHub), 2026-10-19. https://github.com/IABTechLab/seller-agent/tree/feat/openproposal-mapping Accessed 2026-11-10.

[142] First-Party Data Just Got A Lot Smarter: IAS and Databricks CustomerLake. IAS Team (IAS Insider blog), 2026-06-16. https://integralads.com/insider/ias-databricks-customerlake-agentic-cdp-launch-partner/ Accessed 2026-11-10.

[143] Beverage marketer sees cost savings with agentic media-buying test. Sara Karlovitch / Marketing Dive, 2026-03-17 (event 2025-12 to 2026-06). https://www.marketingdive.com/news/beverage-marketer-sees-cost-savings-with-agentic-media-buying-test/814905/ Accessed 2026-11-11.

[144] Agents Ran Entire Butler/Till Campaign, They Share Results. Laurie Sullivan / MediaPost, 2026-03-19 (event 2025-12 to 2026-06). https://www.mediapost.com/publications/article/413607/agents-ran-entire-butlertill-campaign-they-share.html Accessed 2026-11-11.

[145] HUMAN Launches AI-Powered Ad Verification for Advertisers and Agencies Seeking a Modern Alternative. HUMAN Security, Inc. (release distributed by GlobeNewswire), 2026-06-16. https://www.globenewswire.com/news-release/2026/06/16/3312665/0/en/human-launches-ai-powered-ad-verification-for-advertisers-and-agencies-seeking-a-modern-alternative.html Accessed 2026-11-10.

[146] HUMAN Security Unveils AI-Powered Verification Suite for Platforms with OpenX as Launch Partner. HUMAN Security (GlobeNewswire), 2026-10-05. https://www.globenewswire.com/news-release/2026/10/05/3374625/0/en/human-security-unveils-ai-powered-verification-suite-for-platforms-with-openx-as-launch-partner.html Accessed 2026-11-10.

[147] Accreditation letter to Integral Ad Science, 9 October 2024 (GI100924: Recurring and Meta), archived copy. Media Rating Council (MRC), 2024-10-09. https://web.archive.org/web/20250521082231id_/https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/IAS/GI100924_Integral%20Ad%20Science%20Recurring%20and%20Meta%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[148] Accreditation letter to Peer39 for the Peer39 Adloox service (GI062926). Media Rating Council (MRC), 2026-06-29. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Adloox/GI062926_Peer39%20Adloox%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[149] Accreditation letter to Protected Media (GI092226). Media Rating Council (MRC), 2026-09-22. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Protected%20Media/GI092226_Protected_Media_Accreditation_Letter.pdf Accessed 2026-11-10.

[150] Accreditation letter to HUMAN (GI041426: Ad Measurement Platform). Media Rating Council (MRC), 2026-04-14. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/HUMAN/GI041426_HUMAN%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[151] Accreditation letter to Pixalate (GI040825: Analytics platform). Media Rating Council (MRC), 2025-04-08. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Pixalate/GI040825_Pixalate%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[152] Accreditation letter to Zefr: content-level brand safety and suitability on YouTube (GI012026). Media Rating Council (MRC), 2026-01-20. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Zefr/GI012026_Zefr_Google_YouTube_Brand_Safety_Accreditation_Letter.pdf Accessed 2026-11-10.

[153] Accreditation letter to Zefr: third-party reporting of Google ADH YouTube data (GI110625). Media Rating Council (MRC), 2025-11-06. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Zefr/GI110625_Zefr%20Google%20ADH%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[154] ANA Programmatic Media Supply Chain Transparency Study: First Look. Association of National Advertisers (ANA), 2023-06 (event 2023-06-19 (ANA release)). https://info.adtheorent.com/hubfs/Gated-Downloads/ANA%20Programmatic%20Media%20Supply%20Chain%20Transparency%20Study%202023.pdf Accessed 2026-11-10.

[155] DoubleVerify to Protect Advertisers Against Made for Advertising Content with Industry-First Holistic Solution (press release). DoubleVerify, 2023-09-14 (event 2023-09-20 (beta to all advertisers)). https://ir.doubleverify.com/news-events/press-releases/press-releases-details/2023/DoubleVerify-to-Protect-Advertisers-Against-Made-for-Advertising-Content-with-Industry-First-Holistic-Solution-09-14-2023/default.aspx Accessed 2026-11-10.

[156] IAS Launches New Made for Advertising (MFA) AI-Driven Site Detection and Avoidance Technology (press release). Integral Ad Science, 2023-10-16. https://integralads.com/news/ias-new-mfa-site-detection-avoidance-technology/ Accessed 2026-11-10.

[157] Programmatic Transparency Benchmark: 2024 Benchmark, Background Information and Initial Findings. ANA, TAG TrustNet, Fiducia, 2024-06. https://www.fiducia.eco/_files/ugd/de006c_0ff2aed2d03f4b57ac808c037a7f8122.pdf Accessed 2026-11-10.

[158] Programmatic Transparency Benchmark: 2024 Benchmark, Findings Update. ANA, TAG TrustNet, Fiducia, 2024-12 (event 2024-12-12 (ANA release, per the brief-check notes)). https://www.fiducia.eco/_files/ugd/de006c_d67d6426c965479ebb9bcfe6a09548f8.pdf Accessed 2026-11-10.

[159] Programmatic Transparency Benchmark: Q1 2025 Benchmark Findings. ANA, TAG TrustNet, Fiducia, 2025. https://www.fiducia.eco/_files/ugd/de006c_b631a89b157245a2a88e5f22a2cb58f3.pdf Accessed 2026-11-10.

[160] Did Google mislead advertisers about TrueView skippable in-stream ads for the past three years?. Adalytics Research LLC, 2023-06-27. https://adalytics.io/blog/invalid-google-video-partner-trueview-ads Accessed 2026-11-10.

[161] Responses to inquiries with the Media Rating Council on TrueView skippable in-stream ads. Adalytics Research LLC, 2023-07. https://adalytics.io/blog/responses-to-inquiries-with-the-media-rating-council-on-trueview-skippable-in-stream-ads Accessed 2026-11-10.

[162] Transparency and brand safety on Google Video Partners. Google (Marvin Renaud, Director, Global Video Solutions), 2023-06-27 (event updated 2023-07-13). https://blog.google/products/ads-commerce/transparency-and-brand-safety-on-google-video-partners/ Accessed 2026-11-10.

[163] IAS Measurement Across YouTube and Google Video Partners (GVP) Inventory. Integral Ad Science, 2023-07-12. https://integralads.com/apac/insider/ias-measurement-youtube-google-video-partners-inventory/ Accessed 2026-11-10.

[164] Are Amazon and other ad tech vendors serving ANA members' ads on "Made for Arbitrage" websites in 2024?. Adalytics Research LLC, 2024-03. https://adalytics.io/blog/ads-observed-on-made-for-advertising-sites-in-january-2024 Accessed 2026-11-10.

[165] Upholding Accuracy: Addressing Misrepresentations in Recent Third-Party Reports. DoubleVerify, 2024-08-07. https://doubleverify.com/blog/omnichannel/verify/upholding-accuracy-addressing-misrepresentations-in-recent-third-party-reports Accessed 2026-11-10.

[166] Does Forbes operate "made for advertising" pages and transact mis-declared ad inventory?. Adalytics Research LLC, 2024-04-04. https://adalytics.io/blog/ads-observed-on-www3-forbes-subdomain Accessed 2026-11-10.

[167] "AI" brand safety technology. Adalytics Research LLC, 2024-08-07. https://adalytics.io/blog/ai-brand-safety Accessed 2026-11-10.

[168] IAS's Commitment to Media Quality: An Update for Customers and Partners Following a Recent Adalytics Report. Integral Ad Science, 2024-08-07. https://integralads.com/insider/commitment-to-media-quality/ Accessed 2026-11-10.

[169] Are ad tech vendors facilitating or monitoring ads on a website that hosts Child Sexual Abuse Material? (full report; executive summary at /blog/adtech-vendors-csam). Adalytics Research LLC, 2025-02-07. https://adalytics.io/blog/adtech-vendors-csam-full-report Accessed 2026-11-10.

[170] Statement on Adalytics Report. DoubleVerify, 2025-02-07. https://doubleverify.com/blog/web/verify/statement-on-adalytics-report Accessed 2026-11-10.

[171] IAS's Response to the Adalytics Report. Integral Ad Science, 2025-02-07. https://integralads.com/insider/ias-response-adalytics-report/ Accessed 2026-11-10.

[172] Taking Action: Strengthening Protections Against High-Risk, Ad-Supported Content And Attacking the CSAM Issue Together. DoubleVerify, 2025-02-13. https://doubleverify.com/blog/omnichannel/verify/taking-action-strengthening-protections-against-high-risk-ad-supported-content-and-attacking-the-csam-issue-together Accessed 2026-11-10.

[173] On pre-bid bot detection and filtration - Are ad tech vendors serving US Government and Fortune 500 brands' digital ads to bots?. Adalytics Research LLC, 2025-03-28. https://adalytics.io/blog/prebid-bot-filtration Accessed 2026-11-10.

[174] Setting the Record Straight: Addressing Adalytics' Continued Misleading Claims. DoubleVerify, 2025-01-10. https://doubleverify.com/blog/web/verify/setting-the-record-straight-addressing-adalytics-continued-misleading-claims Accessed 2026-11-10.

[175] Statement on Adalytics' GIVT Report. DoubleVerify, 2025-03-28. https://doubleverify.com/blog/web/verify/statement-on-adalytics-givt-report Accessed 2026-11-10.

[176] DoubleVerify's Response to Adalytics' March 28 GIVT Report. DoubleVerify, 2025-04-07. https://doubleverify.com/blog/web/verify/doubleverifys-response-to-adalytics-march-28-givt-report Accessed 2026-11-10.

[177] Update: IAS's Multi-Tiered Approach to Combating Ad Fraud. Integral Ad Science, 2025-03-28. https://integralads.com/insider/update-ias-multi-tiered-approach-to-combatting-ad-fraud/ Accessed 2026-11-10.

[178] Addressing Claims on IAS's Fraud Detection Technology. Integral Ad Science, 2025-04-03. https://integralads.com/insider/addressing-claims-fraud-detection-technology/ Accessed 2026-11-10.

[179] Tens of thousands of news articles are labeled as unsafe for advertisers. Adalytics (Krzysztof Franaszek), 2020-12 (event updated 2021-02-02 (note on the page)). https://adalytics.io/blog/tens-of-thousands-of-news-articles-are-labeled-as-unsafe-for-advertisers Accessed 2026-11-10.

[180] Major brands appear to be blocking ad targeting keywords. Adalytics (Krzysztof Franaszek), 2021-01 (event updated 2021-02-02 (note on the page)). https://adalytics.io/blog/major-brands-appear-to-be-blocking-ad-targeting-keywords Accessed 2026-11-10.

[181] Crossref record for DOI 10.1609/icwsm.v20i1.42744. Crossref (publisher AAAI), 2026-05-25. https://api.crossref.org/works/10.1609/icwsm.v20i1.42744 Accessed 2026-11-10.

[182] Datasheet for Brand Safety Dataset and README, repository rag24/brand_safety_tool_audit. Riley Grossman et al. (NJIT / DeepSee.io), 2026-03. https://github.com/rag24/brand_safety_tool_audit Accessed 2026-11-10.

[183] FP-Inconsistent: Measurement and Analysis of Fingerprint Inconsistencies in Evasive Bot Traffic (arXiv:2406.07647v3). Hari Venugopalan, Shaoor Munir, Shuaib Ahmed, Tangbaihe Wang, Samuel T. King, Zubair Shafiq, 2024-06-11 (event v3 2025-09-21). https://arxiv.org/abs/2406.07647 Accessed 2026-11-10. *Access: abstract only.*

[184] The Inventory is Dark and Full of Misinformation: Understanding the Abuse of Ad Inventory Pooling in the Ad-Tech Supply Chain (arXiv:2210.06654v3). Yash Vekaria, Rishab Nithyanand, Zubair Shafiq, 2022-10-13 (event v3 2023-10-14). https://arxiv.org/abs/2210.06654 Accessed 2026-11-10. *Access: abstract only.*

[185] Building greater transparency and accountability with the Violative View Rate. YouTube (Jennifer O'Connor, Director, YouTube Trust and Safety), 2021-04-06. https://blog.youtube/inside-youtube/building-greater-transparency-and-accountability/ Accessed 2026-11-10.

[186] Prevalence (Meta Transparency Center, "Improving" section). Meta, 2025-03-06. https://transparency.meta.com/policies/improving/prevalence-metric/ Accessed 2026-11-10. *Access: partial extract.*

[187] Media Quality Report: 20th Edition (Americas edition). Integral Ad Science, 2025-05. https://avia.org/wp-content/uploads/2025/05/AMER_MQR_20th_Edition_IAS.pdf Accessed 2026-11-09.

[188] The Economic Cost of Bad Actors on the Internet: Ad Fraud 2020. CHEQ, with Professor Roberto Cavazos (University of Baltimore), 2020. https://info.cheq.ai/hubfs/Research/Economic-Cost-BAD-ACTORS-ON-THE-INTERNET-Ad-Fraud-2020.pdf Accessed 2026-11-09.

[189] Report From ANA And White Ops Shows War On Ad Fraud Is Succeeding (2018-2019 Bot Baseline). Association of National Advertisers (ANA), 2019-05-01. https://ana.net/content/show/id/54065 Accessed 2026-11-09.

[190] 2024 US Ad Fraud Savings Report. TAG, in partnership with the 4A's, ANA and IAB (developed by Scott Cunningham, Cunningham.Tech Consulting), 2024-10-09. https://ppc.land/content/files/2024/10/TAG-US-Fraud-Savings-Report-2024.pdf Accessed 2026-11-09.

[191] Global Study: Fueled by AI, CTV Fraud Schemes Surge 140% Globally (release of DV's 2026 Global Insights report 'Must-CTV: Streaming's Shift From Promise to Performance'). DoubleVerify, 2026-05-07. https://doubleverify.com/company/newsroom/global-study-ctv-fraud-schemes-surge-140-globally Accessed 2026-11-09.

[192] About invalid traffic. Google (Google Ads Help). https://support.google.com/google-ads/answer/11182074?hl=en Accessed 2026-11-10.

[193] Bid on viewable impressions using viewable CPM. Google (Google Ads Help). https://support.google.com/google-ads/answer/3499086?hl=en Accessed 2026-11-10.

[194] Cost-per-thousand viewable impressions (viewable CPM). Google (Google Ads Help). https://support.google.com/google-ads/answer/6026409?hl=en Accessed 2026-11-10.

[195] Resources for Advertisers. Google Ad Traffic Quality. https://www.google.com/ads/adtrafficquality/advertisers/ Accessed 2026-11-10.

[196] Invalid traffic refunds by exchanges. Google (Display & Video 360 Help). https://support.google.com/displayvideo/answer/7556661?hl=en Accessed 2026-11-10.

[197] Filtering invalid traffic to ensure quality. Google (Display & Video 360 Help). https://support.google.com/displayvideo/answer/6076504?hl=en Accessed 2026-11-10.

[198] Position targeting. Google (Display & Video 360 Help). https://support.google.com/displayvideo/answer/2697401?hl=en Accessed 2026-11-10.

[199] Archived feature deprecations (DV360 API). Google for Developers. https://developers.google.com/display-video/api/deprecations-archive Accessed 2026-11-10.

[200] Deductions from earnings FAQs. Google (Google Ad Manager Help). https://support.google.com/admanager/answer/6389769?hl=en Accessed 2026-11-10.

[201] Google Ad Manager 360 billing setup. Google (Google Ad Manager Help). https://support.google.com/admanager/answer/7079586?hl=en Accessed 2026-11-10.

[202] Amazon Advertising Agreement. Amazon (Amazon Ads), 2026-07-23. https://advertising.amazon.com/terms Accessed 2026-11-10.

[203] Traffic validation. Amazon (Amazon Ads Help), 2026-11-02. https://advertising.amazon.com/help/GBYPYH79NGJJ5JPS Accessed 2026-11-10.

[204] How Amazon Ads delivers on quality supply. Amazon (Amazon Ads), 2025-06-27 (event 2026-10-09). https://advertising.amazon.com/library/news/supply-quality Accessed 2026-11-10.

[205] Self-Serve Ad Terms. Meta Platforms, Inc., 2025-08-22. https://www.facebook.com/legal/self_service_ads_terms Accessed 2026-11-10.

[206] Gross impressions. Meta (Business Help Center). https://www.facebook.com/business/help/774724292709131 Accessed 2026-11-10.

[207] TikTok for Business Ad Services Payment Terms. TikTok, 2025-09. https://ads.tiktok.com/i18n/official/policy/payment-terms Accessed 2026-11-10.

[208] TikTok for Business Advertising Terms. TikTok, 2026-01. https://ads.tiktok.com/i18n/official/policy/advertising-terms Accessed 2026-11-10.

[209] TikTok Ad Serving Policy. TikTok, 2024-10. https://ads.tiktok.com/resources/help/article/ad-serving-policy?lang=en Accessed 2026-11-10.

[210] Payment Terms. Snap Inc., 2024-04-01. https://www.snap.com/terms/payment Accessed 2026-11-10.

[211] Self-Serve Advertising Terms. Snap Inc., 2024-09-30. https://www.snap.com/terms/self-serve-advertising Accessed 2026-11-10.

[212] Pinterest Advertising Services Agreement (United States of America). Pinterest, Inc., 2025-04. https://business.pinterest.com/pinterest-advertising-services-agreement/united-states-of-america/ Accessed 2026-11-10.

[213] Reddit Advertising Services Agreement. Reddit, Inc., 2025-12-16 (event 2026-01-01). https://business.reddithelp.com/s/article/Reddit-Advertising-Services-Agreement Accessed 2026-11-10.

[214] Reddit Business Tool Terms. Reddit, Inc., 2025-07-07 (event 2025-08-07). https://business.reddithelp.com/s/article/Reddit-Business-Tool-Terms Accessed 2026-11-10.

[215] X Master Services Agreement - US (with the X Ads Program T&Cs). X Corp., 2026-08-14. https://legal.x.com/ads-terms/us.html Accessed 2026-11-10.

[216] LinkedIn Ads Agreement. LinkedIn Corporation, 2025-11-03. https://www.linkedin.com/legal/sas-terms Accessed 2026-11-10.

[217] OpenPath Publisher Terms and Conditions v.7.5.23. The Trade Desk, Inc., 2023-07. https://assets.thetradedesk.com/assets/documents/legal/URL-OpenPath-Publisher-Terms-and-Conditions-v.7.5.23.pdf Accessed 2026-11-10.

[218] Inventory Supplier Terms and Conditions. The Trade Desk, Inc., 2020-09-24. https://www.thetradedesk.com/legal/inventory-supplier-terms-and-conditions Accessed 2026-11-10.

[219] Trade Desk Advertising Terms. The Trade Desk, Inc., 2026-05-21. https://www.thetradedesk.com/trade-desk-advertising-terms Accessed 2026-11-10.

[220] Microsoft Invest - Understand invalid traffic. Microsoft (Microsoft Learn, Xandr documentation), 2025-04-04. https://learn.microsoft.com/en-us/xandr/invest/understanding-invalid-traffic Accessed 2026-11-10.

[221] Microsoft Invest - Guaranteed Outcomes. Microsoft (Microsoft Learn). https://learn.microsoft.com/en-us/xandr/invest/guaranteed-outcomes Accessed 2026-11-10.

[222] Microsoft Invest - Guaranteed Outcomes mechanism and bid process. Microsoft (Microsoft Learn), 2024-02-15. https://learn.microsoft.com/en-us/xandr/invest/how-guaranteed-outcomes-work Accessed 2026-11-10.

[223] Advertising Terms and Conditions (with DSP, Display Advertising, Creative Services and Foreign Management Program Terms). Yahoo, 2026-01-01. https://legal.yahoo.com/us/en/yahoo/terms/advertising-322/index.html Accessed 2026-11-10.

[224] EXTERNAL: Third Party Vendor List (Yahoo Certification Program). Yahoo, 2026-11-05. https://s.yimg.com/df/3rd_Party_Vendor_List.pdf Accessed 2026-11-10.

[225] Yahoo Exchange Policies. Yahoo (Yahoo Ad Specs). https://adspecs.yahooinc.com/pages/policies-guidelines/yahoo-exchange-policy Accessed 2026-11-10.

[226] Buyer Master Terms and Conditions. Magnite, Inc., 2026-06-17. https://www.magnite.com/legal/buyerterms/ Accessed 2026-11-10.

[227] Inventory Quality Guidelines. Magnite, Inc., 2026-02-09. https://www.magnite.com/legal/inventory-quality-guidelines/ Accessed 2026-11-10.

[228] Demand Partner Master Services Agreement. PubMatic, Inc., 2021-05-28 (event 2026-08-18). https://pubmatic.com/legal/demand-service-agreement/ Accessed 2026-11-10.

[229] Publisher Master Services Agreement. PubMatic, Inc., 2019-01-03 (event 2025-09-05). https://pubmatic.com/legal/publisher-master-services-agreement/ Accessed 2026-11-10.

[230] Supply Policy. PubMatic, Inc., 2026-10-07. https://pubmatic.com/legal/supply-policy/ Accessed 2026-11-10.

[231] Digital Ad Inventory Quality. PubMatic, Inc.. https://pubmatic.com/inventory-quality/ Accessed 2026-11-10.

[232] Master Media Buying Services Agreement Template, Version 3.0. Association of National Advertisers (ANA), 2023-06. https://www.ana.net/getfile/35568 Accessed 2026-11-10.

[233] ANA Releases Updated Media Buying Contract Template. Association of National Advertisers (ANA), 2023-06-06. https://www.ana.net/content/show/id/pr-2023-06-updated-media-buying-contract Accessed 2026-11-10.

[234] General Terms and Conditions for Digital Advertising Services, Version 1.1. Interactive Advertising Bureau (IAB), 2026-05. https://www.iab.com/wp-content/uploads/2026/05/IAB-General-Terms-v1.1-Final_May-2026.pdf Accessed 2026-11-10.

[235] Direct Buy Addendum to the General Terms and Conditions for Digital Advertising Services, Version 1.0. Interactive Advertising Bureau (IAB), 2026-05. https://www.iab.com/wp-content/uploads/2026/05/IAB-Direct-Buy-Addendum-to-General-Terms-and-Conditions-v1.0-Final_May-2026.pdf Accessed 2026-11-10.

[236] Appendix 1 (Definitions) to the General Terms and Conditions for Digital Advertising Services and Specific Terms, Version 1.1. Interactive Advertising Bureau (IAB), 2026-05. https://www.iab.com/wp-content/uploads/2026/05/Appendix-1-Definitions-v1.1-Final_May-2026.pdf Accessed 2026-11-10.

[237] Measurement Services Addendum v1.0 Final. Interactive Advertising Bureau (IAB), 2026-10-28. https://www.iab.com/guidelines/iab-releases-measurement-services-addendum-for-public-comment/ Accessed 2026-11-10. *Access: partial extract.*

[238] Measurement Services Addendum to the General Terms and Conditions for Digital Advertising Services, Version 1.0 (file). Interactive Advertising Bureau (IAB), 2026-10. https://www.iab.com/wp-content/uploads/2026/10/IAB_Measurement_Addendum_v1.0_Final_Oct_2026.docx Accessed 2026-11-10.

[239] 4A's/IAB Standard Terms and Conditions for Internet Advertising for Media Buys One Year or Less, Version 3.0. American Association of Advertising Agencies (4A's) and Interactive Advertising Bureau (IAB), 2009. https://www.iab.com/wp-content/uploads/2015/06/IAB_4As-tsandcs-FINAL.pdf Accessed 2026-11-10.

[240] DAS IAB Addendum. Disney Advertising, 2023-03-28. https://www.disneyadvertising.com/mediakit/das-iab-addendum/ Accessed 2026-11-10.

[241] Spotify Master Advertising Terms and Conditions. Spotify, 2026-07-01. https://www.spotify.com/us/brands/legal/advertiser-terms-and-conditions/ Accessed 2026-11-10.

[242] Content Delivery Reports API (Marketing API, Brand Safety and Suitability). Meta for Developers. https://developers.facebook.com/documentation/ads-commerce/marketing-api/brand-safety-and-suitability/content-delivery-report Accessed 2026-11-10.

[243] Third-party measurement on YouTube (Google Ads Help). Google. https://support.google.com/google-ads/answer/9741375?hl=en Accessed 2026-11-10.

[244] Accreditation letter to Google: Ads Data Hub platform. Media Rating Council (George W. Ivie), 2026-01-20. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Google/GI012026%20Google%20Ads%20Data%20Hub%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[245] Frequently asked questions (Google Ads Help; YouTube Shorts). Google. https://support.google.com/google-ads/answer/16041383?hl=en Accessed 2026-11-10.

[246] Zefr Launches Brand Safety & Suitability Reporting (BSRP) on YouTube, Providing Advertisers with Third-Party Measurement. Zefr, 2022-10-31. https://zefr.com/press/zefr-launches-brand-safety-suitability-reporting-bsrp-on-youtube-providing-advertisers-with-third-party-measurement Accessed 2026-11-10.

[247] DoubleVerify Partners with TikTok to Measure Viewability and Invalid Traffic (Business Wire release, syndicated copy). DoubleVerify (distributed by Business Wire; copy hosted on investor.wedbush.com), 2021-09-28. https://investor.wedbush.com/wedbush/article/bizwire-2021-9-28-doubleverify-partners-with-tiktok-to-measure-viewability-and-invalid-traffic Accessed 2026-11-10.

[248] IAS Expands Partnership with TikTok to Measure Viewability and Invalid Traffic Globally. Integral Ad Science, 2022-05-09. https://integralads.com/news/ias-expands-partnership-with-tiktok-to-measure-viewability-and-invalid-traffic-globally/ Accessed 2026-11-10.

[249] Open Measurement (OM) SDK. IAB Tech Lab, 2025-11-04. https://iabtechlab.com/standards/open-measurement-sdk/ Accessed 2026-11-10.

[250] IAS Expands Partnership with TikTok to Provide Leading Comprehensive Third-Party Brand Safety Measurement Suite for the Platform. Integral Ad Science, 2022-10-17. https://integralads.com/uk/news/ias-expands-partnership-with-tiktok-to-provide-leading-comprehensive-third-party-brand-safety-measurement-suite-for-the-platform/ Accessed 2026-11-10.

[251] Introducing TikTok's Brand Safety and Suitability Partners. TikTok for Business (blog), 2023-05-19. https://ads.tiktok.com/business/en/blog/introducing-brand-safety-suitability-partners Accessed 2026-11-10.

[252] Approved third-party providers (Amazon Ads ad policy). Amazon Ads, 2026-06-01. https://advertising.amazon.com/resources/ad-policy/en/approved-3p-ad-servers Accessed 2026-11-10.

[253] DoubleVerify Expands Global Measurement Capabilities To Include Amazon Custom Audiences. DoubleVerify, 2023-10-31. https://doubleverify.com/company/newsroom/doubleverify-expands-global-measurement-capabilities-to-include-amazon-custom-audiences Accessed 2026-11-10.

[254] DoubleVerify Debuts Industry-Leading In-Feed Brand Safety and Suitability Solution on Twitter (Business Wire release, DV investor-relations PDF). DoubleVerify, 2023-01-25. https://s206.q4cdn.com/961864615/files/doc_news/2023/01/2023-01-25_DoubleVerify_Debuts_Industry_Leading_In_Feed_223.pdf Accessed 2026-11-10.

[255] Types of Third-Party Tags and Supported Actions (Snapchat Business Help Center). Snap Inc.. https://businesshelp.snapchat.com/s/article/third-party-tag-types?language=en_US Accessed 2026-11-10.

[256] Third Party Tracking Overview (Tag-Based Solutions) (Snapchat Business Help Center). Snap Inc.. https://businesshelp.snapchat.com/s/article/third-party-tracking?language=en_US Accessed 2026-11-10.

[257] Ad Squads (Snap Marketing API). Snap Inc. (Snap for Developers). https://developers.snap.com/marketing-api/Ads-API/ad-squads Accessed 2026-11-10.

[258] IAS Brand Safety Reporting (Snapchat Business Help Center). Snap Inc.. https://businesshelp.snapchat.com/s/article/ias-brand-safety-reporting?language=en_US Accessed 2026-11-10.

[259] DoubleVerify Bolsters Media Quality Authentication on Pinterest with Brand Safety and Suitability Measurement. DoubleVerify, 2024-06-13. https://doubleverify.com/company/newsroom/doubleverify-bolsters-media-quality-authentication-on-pinterest-with-brand-safety-and-suitability-measurement Accessed 2026-11-10.

[260] IAS Now Offers Fraud and Viewability Measurement Across Pinterest. Integral Ad Science, 2020-08-19. https://integralads.com/news/pinterest-fraud-viewability-measurement/ Accessed 2026-11-10.

[261] Brand Safety on Pinterest. Pinterest (Pinterest Business). https://business.pinterest.com/brand-safety/ Accessed 2026-11-10.

[262] DoubleVerify Unlocks Comprehensive Media Quality Authentication on Reddit with First-to-Market Offering. DoubleVerify, 2024-06-11. https://doubleverify.com/company/newsroom/doubleverify-unlocks-comprehensive-media-quality-authentication-on-reddit-with-first-to-market-offering Accessed 2026-11-10.

[263] IAS Expands Partnership With Reddit to Provide Additional Third-Party Measurement for Advertisers. Integral Ad Science, 2025-04-03. https://integralads.com/news/reddit-partnership-expansion-measurement-for-advertisers/ Accessed 2026-11-10.

[264] Click and Impression Trackers (Reddit Ads Help). Reddit, Inc.. https://business.reddithelp.com/s/article/Set-up-third-party-measurement Accessed 2026-11-10.

[265] Ad Inventory Types (Reddit Ads Help). Reddit, Inc.. https://business.reddithelp.com/s/article/About-ad-inventory-types Accessed 2026-11-10.

[266] Get started with third-party tracking for LinkedIn Ads (LinkedIn Marketing Solutions Help). LinkedIn (Microsoft). https://www.linkedin.com/help/lms/answer/a420844 Accessed 2026-11-10.

[267] Set up DoubleVerify measurement for your LinkedIn Ads (LinkedIn Marketing Solutions Help). LinkedIn (Microsoft). https://www.linkedin.com/help/lms/answer/a10959097 Accessed 2026-11-10.

[268] Set up Integral Ad Science (IAS) measurement for your video ads (LinkedIn Marketing Solutions Help). LinkedIn (Microsoft). https://www.linkedin.com/help/lms/answer/a16384016 Accessed 2026-11-10.

[269] Brand Safety & Suitability Hub in Campaign Manager (LinkedIn Help). LinkedIn (Microsoft). https://www.linkedin.com/help/lms/answer/a1337901 Accessed 2026-11-10.

[270] Quality and brand safety for the LinkedIn Audience Network (LinkedIn Help). LinkedIn (Microsoft). https://www.linkedin.com/help/linkedin/answer/a422485 Accessed 2026-11-10.

[271] DoubleVerify Holdings, Inc. Form 10-K for the fiscal year ended 31 December 2024. DoubleVerify Holdings, Inc. (SEC filing), 2025-02-27 (event 2024-12-31). https://www.sec.gov/Archives/edgar/data/1819928/000155837025001814/dv-20241231×10k.htm Accessed 2026-11-09.

[272] DoubleVerify Holdings, Inc. Form 10-K for the fiscal year ended 31 December 2023. DoubleVerify Holdings, Inc. (SEC filing), 2024-02-28 (event 2023-12-31). https://www.sec.gov/Archives/edgar/data/1819928/000155837024002008/dv-20231231×10k.htm Accessed 2026-11-09.

[273] DoubleVerify Holdings, Inc. Form 10-Q for the quarter ended 30 September 2022. DoubleVerify Holdings, Inc. (SEC filing), 2022-11-08 (event 2022-09-30). https://www.sec.gov/Archives/edgar/data/1819928/000155837022016885/dv-20220930×10q.htm Accessed 2026-11-09.

[274] Nielsen to Acquire DoubleVerify, Creating a Leading, Independent Media Intelligence Platform (joint release, Exhibit 99.1 to DoubleVerify's 8-K of 6 August 2026). Nielsen and DoubleVerify (joint press release filed by DoubleVerify), 2026-08-06. https://www.sec.gov/Archives/edgar/data/1819928/000110465926092121/tm2621904d1_ex99-1.htm Accessed 2026-11-09.

[275] DoubleVerify Holdings, Inc. Form 8-K (Item 5.07, special meeting results). DoubleVerify Holdings, Inc. (SEC filing), 2026-10-29. https://www.sec.gov/Archives/edgar/data/1819928/000110465926120260/tm2627835d3_8k.htm Accessed 2026-11-09.

[276] Integral Ad Science Holding Corp. Form 10-K for the fiscal year ended 31 December 2021. Integral Ad Science Holding Corp. (SEC filing), 2022-03-03 (event 2021-12-31). https://www.sec.gov/Archives/edgar/data/1842718/000162828022004899/ias-20211231.htm Accessed 2026-11-09.

[277] Integral Ad Science Holding Corp. Form 8-K (Items 1.01, 5.07, 9.01: Merger Agreement). Integral Ad Science Holding Corp. (SEC filing), 2025-09-24. https://www.sec.gov/Archives/edgar/data/1842718/000119312525215824/d26134d8k.htm Accessed 2026-11-10.

[278] IAS to be Acquired by Novacap for $1.9 Billion to Further Support Its Strategic Goals (joint release, Exhibit 99.1 to 8-K of 24 September 2025). Integral Ad Science and Novacap (joint press release filed by IAS), 2025-09-24. https://www.sec.gov/Archives/edgar/data/1842718/000119312525214050/d938971dex991.htm Accessed 2026-11-09.

[279] White Ops Reintroduces Itself as HUMAN, Signaling Cybersecurity Company's Dedication to Protecting Enterprises from Bot Attacks. HUMAN Security, 2021-03-30. https://web.archive.org/web/20210330125918id_/https://www.humansecurity.com/newsroom/white-ops-reintroduces-itself-as-human-signaling-cybersecurity-companys-dedication-to-protecting-enterprises-from-bot-attacks Accessed 2026-11-09.

[280] HUMAN and PerimeterX unite in market-changing merger to safeguard customers from sophisticated bot attacks, fraud and account abuse. HUMAN Security and PerimeterX, 2022-07-27. https://web.archive.org/web/20220727122104id_/https://www.humansecurity.com/newsroom/human-and-perimeterx-unite-in-market-changing-merger-to-safeguard-customers-from-sophisticated-bot-attacks-fraud-and-account-abuse Accessed 2026-11-09.

[281] Adloox is now part of Scope3. Adloox, 2024-11-05 (event 2024-11-04). https://www.adloox.com/blog/adloox-is-now-part-of-scope3 Accessed 2026-11-09.

[282] Peer39 Acquires Adloox from Scope3, Unifying Contextual Intelligence, Brand Safety Across Google, Meta, and the Open Web. Peer39 (release distributed by EIN Presswire), 2026-06-02. https://newsroom.xpr.media/story/155067/peer39-acquires-adloox-from-scope3-unifying-contextual-intelligence-brand-safety-across-google-meta-and-the-open-web/ Accessed 2026-11-09.

[283] GARM's Harm: How the World's Biggest Brands Seek to Control Online Speech (Interim Staff Report). Committee on the Judiciary, U.S. House of Representatives (majority staff), 2024-07-10. https://judiciary.house.gov/sites/evo-subsites/republicans-judiciary.house.gov/files/evo-media-document/2024-07-10%20GARMs%20Harm%20-%20How%20the%20Worlds%20Biggest%20Brands%20Seek%20to%20Control%20Online%20Speech.pdf Accessed 2026-11-10.

[284] X Corp. v. World Federation of Advertisers et al., Memorandum Opinion and Order (ECF 255), 7:24-cv-00114-B (N.D. Tex.). U.S. District Court, N.D. Tex. (Senior Judge Jane J. Boyle), 2026-03-26. https://storage.courtlistener.com/recap/gov.uscourts.txnd.393003/gov.uscourts.txnd.393003.255.0.pdf Accessed 2026-11-10.

[285] GARM (WFA page carrying the discontinuation statement). World Federation of Advertisers (WFA), 2024-08-09. https://wfanet.org/leadership/garm/about-garm Accessed 2026-11-10.

[286] DOJ, NCIS ask ad executives about brand-safety companies. Marketing Brew (Ryan Barwick), 2024-10-11 (event 2024). https://www.marketingbrew.com/stories/2024/10/11/doj-ncis-brand-safety-google-integral-ad-science-doubleverify Accessed 2026-11-10. *Access: partial extract.*

[287] Oklahoma Firefighters v. IAS, docket 1:25-cv-00847 (S.D.N.Y.). CourtListener (Free Law Project), 2025-01-29 (event 2026-03-26). https://www.courtlistener.com/docket/69588890/oklahoma-firefighters-pension-and-retirement-system-v-integral-ad-science/ Accessed 2026-11-10.

[288] Tauber v. Utzschneider (In re Integral Ad Science Holding Corp. Derivative Litigation), Verified Shareholder Derivative Complaint (ECF 1), 1:25-cv-02039 (S.D.N.Y.). Harvey Tauber (plaintiff), 2025-03-11. https://storage.courtlistener.com/recap/gov.uscourts.nysd.638442/gov.uscourts.nysd.638442.1.0.pdf Accessed 2026-11-10.

[289] Martin v. Zagorski, 1:25-cv-02715 (S.D.N.Y.), docket. CourtListener (Free Law Project), from PACER/RECAP, 2025-03-31. https://www.courtlistener.com/docket/69831835/martin-v-zagorski/ Accessed 2026-11-10.

[290] Pixalate, Inc. v. Criteo S.A., docket 1:25-cv-03614 (S.D.N.Y.). CourtListener (Free Law Project), 2025-04-30 (event 2025-06-18). https://www.courtlistener.com/docket/70013065/pixalate-inc-v-criteo-sa/ Accessed 2026-11-10. *Access: partial extract.*

[291] DoubleVerify Holdings Inc. v. Adalytics Research, LLC, Complaint (ECF 1), 8:25-cv-01535 (D. Md.). DoubleVerify Holdings Inc. and DoubleVerify Inc. (plaintiffs), 2025-05-12. https://storage.courtlistener.com/recap/gov.uscourts.mdd.582283/gov.uscourts.mdd.582283.1.0.pdf Accessed 2026-11-10.

[292] FTC Civil Investigative Demand to The Global Disinformation Index, FTC File No. 251-0061 (Exhibit B to ECF 29, Disinformation Index, Inc. v. FTC, 1:25-cv-04137 (D.D.C.)). U.S. Federal Trade Commission (as filed by GDI), 2025-05-20. https://storage.courtlistener.com/recap/gov.uscourts.dcd.287176/gov.uscourts.dcd.287176.29.2.pdf Accessed 2026-11-10.

[293] Rumble Inc. et al. v. World Federation of Advertisers et al., Memorandum Opinion and Order (ECF 69), 7:24-cv-0115-B (N.D. Tex.). U.S. District Court, N.D. Tex. (Judge Jane J. Boyle), 2025-08-13. https://cases.justia.com/federal/district-courts/texas/txndce/7:2024cv00115/393019/69/0.pdf Accessed 2026-11-10.

[294] Integral Ad Science Holding Corp. Form 10-Q for the quarter ended 30 September 2025. Integral Ad Science Holding Corp. (SEC filing), 2025-11-04 (event 2025-09-30). https://www.sec.gov/Archives/edgar/data/1842718/000184271825000109/ias-20250930.htm Accessed 2026-11-09.

[295] Kaszirer v. Zagorski, Verified Shareholder Derivative Complaint (ECF 1), 1:25-cv-10200 (S.D.N.Y.). Susana Kaszirer (plaintiff), counsel Joshua M. Lifshitz, 2025-12-09. https://storage.courtlistener.com/recap/gov.uscourts.nysd.654410/gov.uscourts.nysd.654410.1.0.pdf Accessed 2026-11-10.

[296] IAS Enhancements to Context Control Avoidance (news release). Integral Ad Science, 2025-12-18. https://integralads.com/news/ias-enhancements-to-context-control-avoidance/ Accessed 2026-11-10.

[297] In re DoubleVerify Holdings, Inc. Securities Litigation, Notice of Voluntary Dismissal (ECF 35) and docket. Lead plaintiffs (Labaton Keller Sucharow LLP); CourtListener docket 70342418, 2025-12-22. https://storage.courtlistener.com/recap/gov.uscourts.nysd.643090/gov.uscourts.nysd.643090.35.0.pdf Accessed 2026-11-10.

[298] In re Integral Ad Science Holding Corp. Derivative Litigation, docket (1:25-cv-02039). CourtListener (Free Law Project), 2025-03-11 (event 2026-01-15). https://www.courtlistener.com/docket/69726541/in-re-integral-ad-science-holding-corp-derivative-litigation/ Accessed 2026-11-10. *Access: partial extract.*

[299] Disinformation Index, Inc. v. FTC, First Amended Complaint (ECF 29). Disinformation Index, Inc. (GDI), plaintiff, 2026-02-02. https://storage.courtlistener.com/recap/gov.uscourts.dcd.287176/gov.uscourts.dcd.287176.29.0.pdf Accessed 2026-11-10.

[300] Kaszirer v. Zagorski, docket and entry pages for ECF 4 and 5. CourtListener (Free Law Project), 2025-12-09 (event 2026-02-06). https://www.courtlistener.com/docket/72011643/kaszirer-v-zagorski/ Accessed 2026-11-10. *Access: partial extract.*

[301] comScore, Inc., Form 10-K for fiscal year 2025, Note 11 Commitments and Contingencies (Privacy Litigation). comScore, Inc. (SEC filing), 2026-03-26. https://www.sec.gov/Archives/edgar/data/1158172/000115817226000009/scor-20251231.htm Accessed 2026-11-10.

[302] Oklahoma Firefighters Pension and Retirement System v. Integral Ad Science Holding Corp., Memorandum Opinion Granting Motions to Dismiss (ECF 62), 25-cv-0847 (LAK). U.S. District Court, S.D.N.Y. (Judge Lewis A. Kaplan), 2026-03-02. https://storage.courtlistener.com/recap/gov.uscourts.nysd.635896/gov.uscourts.nysd.635896.62.0.pdf Accessed 2026-11-10.

[303] FTC et al. v. Dentsu US, Inc., Final Order and Stipulated Permanent Injunction as to Dentsu US, Inc. (ECF 8). U.S. District Court, N.D. Tex.; FTC and eight states; Dentsu US, Inc., 2026-04-15. https://storage.courtlistener.com/recap/gov.uscourts.txnd.418330/gov.uscourts.txnd.418330.8.0.pdf Accessed 2026-11-10.

[304] Disinformation Index, Inc. v. FTC, Notice of Withdrawal of Civil Investigative Demand (ECF 37) and Exhibit 1 (FTC letter of 16 Apr 2026). U.S. Department of Justice for the FTC; FTC Bureau of Competition, 2026-04-21 (event 2026-04-16). https://storage.courtlistener.com/recap/gov.uscourts.dcd.287176/gov.uscourts.dcd.287176.37.0.pdf Accessed 2026-11-10.

[305] DoubleVerify Holdings Inc. v. Adalytics Research, LLC, Memorandum Opinion (ECF 30). U.S. District Court, D. Md. (Judge Theodore D. Chuang), 2026-04-27. https://storage.courtlistener.com/recap/gov.uscourts.mdd.582283/gov.uscourts.mdd.582283.30.0_1.pdf Accessed 2026-11-10.

[306] Chavez v. Human Security, Inc., Class Action Complaint (ECF 1), 1:26-cv-03727 (S.D.N.Y.). Olga Chavez and Andrew Garcia (plaintiffs), 2026-05-05. https://storage.courtlistener.com/recap/gov.uscourts.nysd.663381/gov.uscourts.nysd.663381.1.0.pdf Accessed 2026-11-10.

[307] DoubleVerify v. Adalytics, Answer to Original Complaint (ECF 32). Adalytics Research, LLC (defendant), 2026-05-11. https://storage.courtlistener.com/recap/gov.uscourts.mdd.582283/gov.uscourts.mdd.582283.32.0.pdf Accessed 2026-11-10.

[308] Chavez v. Human Security Inc., docket. CourtListener (Free Law Project), 2026-05-05 (event 2026-07-21). https://www.courtlistener.com/docket/73300644/chavez-v-human-security-inc/ Accessed 2026-11-10.

[309] WFA and X settle over GARM litigation. WFA, 2026-07-29. https://wfanet.org/knowledge/item/2026/07/29/wfa-and-x-settle-over-garm-litigation Accessed 2026-11-10.

[310] Brief for the United States of America as Amicus Curiae in Support of Plaintiff-Appellant (Doc. 126), X Corp. v. Mars, Inc., 5th Cir. No. 26-10394. U.S. Department of Justice (Associate Attorney General; Antitrust Division), 2026-08-12. https://www.justice.gov/atr/media/1458686/dl?inline= Accessed 2026-11-10.

[311] Brief of Amici Curiae State of West Virginia and 11 Other States in Support of Plaintiff-Appellant/Cross-Appellee & Reversal (Doc. 123). Office of the West Virginia Attorney General and 11 states, 2026-08-12. https://business.cch.com/ald/xVmars-123-Statesbrief081826.pdf Accessed 2026-11-10.

[312] DoubleVerify v. Adalytics, docket (8:25-cv-01535). CourtListener (Free Law Project), 2025-05-12 (event 2026-09-02). https://www.courtlistener.com/docket/70251967/doubleverify-holdings-inc-v-adalytics-research-llc/ Accessed 2026-11-10.

[313] Martin v. Zagorski, No. 25-cv-2715 (LJL), Opinion and Order (ECF 61). U.S. District Court, S.D.N.Y. (Judge Lewis J. Liman), 2026-09-16. https://storage.courtlistener.com/recap/gov.uscourts.nysd.639801/gov.uscourts.nysd.639801.61.0.pdf Accessed 2026-11-10.

[314] Michael Kent v. DoubleVerify Holdings, Inc. et al (New York County Supreme Court), case page. UniCourt Inc. (aggregator of state court records), 2026-10-08. https://unicourt.com/case/ny-sue1-casegg62b36ac7777e-3691660 Accessed 2026-11-10. *Access: partial extract.*

[315] ACCC Acquisitions Register: Nielsen - DoubleVerify (MN-55036). Australian Competition and Consumer Commission, 2026-10-13. https://www.accc.gov.au/public-registers/acquisitions-and-mergers-registers/acquisitions-register/nielsen-doubleverify Accessed 2026-11-10.

[316] DoubleVerify Holdings, Inc. Form 8-K (Item 8.01, supplemental proxy disclosures and regulatory status). DoubleVerify Holdings, Inc. (SEC filing), 2026-10-20. https://www.sec.gov/Archives/edgar/data/1819928/000110465926117471/tm2627835d1_8k.htm Accessed 2026-11-09.

[317] DoubleVerify Holdings, Inc. Form 10-Q for the quarter ended 30 September 2026. DoubleVerify Holdings, Inc. (SEC filing), 2026-10-30 (event 2026-09-30). https://www.sec.gov/Archives/edgar/data/1819928/000110465926120666/dv-20260930×10q.htm Accessed 2026-11-09.

[318] Measurement Taxonomy (AdCP documentation). AgenticAdvertising.org (AdCP), 2026-05-02 (event 2026-10-29). https://docs.adcontextprotocol.org/docs/measurement/taxonomy Accessed 2026-11-10.

[319] AdCP JSON schemas at tag v3.2.9: core/performance-standard.json, enums/performance-standard-metric.json, core/measurement-terms.json, enums/makegood-remedy.json. AgenticAdvertising.org (AdCP) on GitHub, 2026-11-09. https://github.com/adcontextprotocol/adcp/tree/v3.2.9/static/schemas/source Accessed 2026-11-10.

[320] AdCP JSON schemas at tags v3.0.0 and v3.1.0 (performance-standard-metric.json, measurement-terms.json, makegood-remedy.json, protocol/get-adcp-capabilities-response.json). AgenticAdvertising.org (AdCP) on GitHub, 2026-04-22 (event 2026-06-18). https://github.com/adcontextprotocol/adcp/tree/v3.0.0/static/schemas/source Accessed 2026-11-10.

[321] Apostra developer documentation corpus (llms-full.txt), including the API v2 changelog. Apostra, Inc. (formerly Scope3), 2026-11-10. https://docs.apostra.com/llms-full.txt Accessed 2026-11-10.

[322] AdCP compliance storyboards at v3.2.9 (static/compliance/source: specialisms/content-standards, protocols/media-buy/scenarios measurement_terms_accepted, billing_finality_delivery, vendor_metric_accountability). AgenticAdvertising.org (AdCP), 2026-11-09. https://github.com/adcontextprotocol/adcp/tree/v3.2.9/static/compliance/source Accessed 2026-11-10.

[323] From Brief to Buy: How Butler/Till Used DV Neura to Embed Verification Into Agentic Workflows. Ben Cann, SVP Product Management, DoubleVerify (blog), 2026-07-24. https://doubleverify.com/blog/omnichannel/optimize/from-brief-to-buy-how-butler/till-used-dv-neura-to-embed-verification-into-agentic-workflows Accessed 2026-11-10.

[324] Content Standards protocol: overview, implementation guide and schemas (validate-content-delivery-response.json, enums/binary-verdict.json). AgenticAdvertising.org (AdCP), 2026-03-01. https://github.com/adcontextprotocol/adcp/tree/v3.2.9/docs/governance/content-standards Accessed 2026-11-10.

[325] AgenticAdvertising.org registry API: agents with capabilities (tool lists as crawled by AAO). AgenticAdvertising.org, 2026-11-10. https://agenticadvertising.org/api/registry/agents?capabilities=true Accessed 2026-11-10.

[326] AgenticAdvertising.org registry API: public compliance records for the 35 listed agents. AgenticAdvertising.org, 2026-11-10. https://agenticadvertising.org/api/registry/agents/{encoded agent URL}/compliance Accessed 2026-11-10.

[327] OpenProposal - Object Model Specification (AAMP 3.0), spec/openproposal-3.0.md and examples/line-item-display.yaml. IAB Tech Lab (GitHub IABTechLab/OpenProposal), 2026-09-22. https://github.com/IABTechLab/OpenProposal/blob/main/spec/openproposal-3.0.md Accessed 2026-11-10.

[328] IAB Tech Lab Introduces AAMP 3.0 with OpenProposal (press release). IAB Tech Lab, 2026-09-22 (event 2026-10-23). https://iabtechlab.com/press-releases/iab-tech-lab-introduces-aamp-3-0-with-openproposal/ Accessed 2026-11-10.

[329] AAMP buyer agent reference implementation (IABTechLab/buyer-agent, ad-buyer-system 2.4.2). IAB Tech Lab (GitHub), 2026-11-06. https://github.com/IABTechLab/buyer-agent Accessed 2026-11-10.

[330] AAMP seller agent (IABTechLab/seller-agent 2.4.2) and iab-agentic-primitives (spec/jsonschema/Deal.json). IAB Tech Lab (GitHub), 2026-11-10. https://github.com/IABTechLab/iab-agentic-primitives Accessed 2026-11-10.

[331] Agentic Real Time Framework, Version 1.0 (FINAL), with docs/intents/add-metrics.md and README. IAB Tech Lab (GitHub IABTechLab/agentic-real-time-framework), 2025-11-12. https://github.com/IABTechLab/agentic-real-time-framework/blob/main/Agentic_Real_Time_Framework_Version_1_0_FINAL.md Accessed 2026-11-10.

[332] Agentic Advertising Matures with AAMP 2.3. IAB Tech Lab, 2026-07-30 (event 2026-08-08). https://iabtechlab.com/agentic-advertising-matures-with-aamp-2-3/ Accessed 2026-11-10.

[333] DoubleVerify Introduces DV Neura, the Dynamic AI Engine Powering the Future of Media Quality and Effectiveness. DoubleVerify (newsroom), 2026-06-17. https://doubleverify.com/company/newsroom/doubleverify-introduces-dv-neura-the-dynamic-ai-engine-powering-the-future-of-media-quality-and-effectiveness Accessed 2026-11-10.

[334] DV Neura product page. DoubleVerify. https://doubleverify.com/products/neura Accessed 2026-11-10.

[335] DV Launches Ask Neura, Expanding Natural-Language Access to Media Intelligence. DoubleVerify (newsroom), 2026-09-17. https://doubleverify.com/company/newsroom/dv-launches-ask-neura-expanding-natural-language-access-to-media-intelligence Accessed 2026-11-10.

[336] IAS Launches 'IAS Agent' Enabling Marketers to Surface Deep Campaign Insights in Real-Time and Increase Ad Performance. Integral Ad Science (newsroom), 2025-12-16. https://integralads.com/news/ias-launches-ias-agent/ Accessed 2026-11-10.

[337] IAS Agent: Always-On Campaign Insights Built for Action. IAS Team (IAS Insider blog), 2026-06-05. https://integralads.com/insider/ias-agent-always-on-campaign-insights-built-for-action/ Accessed 2026-11-10.

[338] IAS Agent product page. Integral Ad Science, 2026-01-22 (event 2026-04-21). https://integralads.com/ias-agent/ Accessed 2026-11-10.

[339] HTTP checks of /.well-known/brand.json on twelve domains (this stream's log). No Fluff Advisory research log (this stream), 2026-11-10. https://doubleverify.com/.well-known/brand.json Accessed 2026-11-10.

[340] Talk to Your Threat Data: HUMAN Security MCP Server for AI-Powered Security Intelligence. Ben Diamant and Ori Gold, HUMAN Security (blog), 2025-07-24 (event 2026-01-12). https://www.humansecurity.com/learn/blog/connect-ai-to-security-data-with-human-mcp-server/ Accessed 2026-11-10.

[341] Introducing HUMAN Verified AI Agent: An Open-Source Foundation for Trustworthy Agent Identity. Ben Diamant and Tomer Elias, HUMAN Security (blog), 2025-07-21 (event 2026-02-27). https://www.humansecurity.com/learn/blog/human-verified-ai-agent-open-source/ Accessed 2026-11-10.

[342] AgenticAdvertising.org public member profiles (members API). AgenticAdvertising.org (AAO), live list. https://agenticadvertising.org/api/members?limit=500 Accessed 2026-11-10.

[343] Zefr Introduces AdCP-Based Agentic Hub for Cross-Platform Social Campaigns. Zefr (press page), 2026-05-01 (event 2026-05-04). https://zefr.com/press/zefr-introduces-adcp-based-agentic-hub-for-cross-platform-social-campaigns Accessed 2026-11-10.

[344] Zefr Expands Its Agentic Infrastructure to Cover the Full Spectrum of YouTube Video Buying. Zefr (press page), 2026-09-14 (event 2026-09-15). https://zefr.com/press/zefr-expands-its-agentic-infrastructure-to-cover-the-full-spectrum-of-youtube-video-buying Accessed 2026-11-10.

[345] Methodology (Brand Standards guide, Scope3 developer docs) and docs.scope3.com/llms.txt. Scope3 (now Apostra), 2025-08-28. https://docs.scope3.com/docs/classification-methodology.md Accessed 2026-11-10.

[346] Apostra brand.json (/.well-known/brand.json). Apostra, Inc. (formerly Scope3), 2026-10-01. https://apostra.com/.well-known/brand.json Accessed 2026-11-10.

[347] Pixalate documentation: MCPs and agents; IQ Agent. Pixalate, Inc., undated (live page). https://docs.pixalate.com/mcp/overview Accessed 2026-11-10.

[348] DoubleVerify Introduces DV AI Verification Offering to Identify and Manage Agent Interactions and Avoid AI Slop. DoubleVerify (newsroom), 2025-11-04. https://doubleverify.com/company/newsroom/doubleverify-introduces-dv-ai-verification-offering-to-identify-and-manage-agent-interactions-and-avoid-ai-slop Accessed 2026-11-10.

[349] PubMatic and Butler/Till Launch AdCP-Enabled Agentic AI Campaign Across CTV. PubMatic (news), 2025-12-15. https://pubmatic.com/news/pubmatic-and-butler-till-launch-adcp-enabled-agentic-ai-campaign-across-

ctv/ Accessed 2026-11-10.

[350] GitHub API: adcontextprotocol/adcp releases (124 listed) and commit histories of three documentation files. GitHub (repository metadata), 2026-11-10. https://api.github.com/repos/adcontextprotocol/adcp/releases Accessed 2026-11-10.

[351] PubMatic Launches Governance Architecture for Agentic Advertising on AgenticOS. PubMatic (news), 2026-08-05. https://pubmatic.com/news/pubmatic-launches-governance-architecture-for-agentic-advertising-on-agenticos/ Accessed 2026-11-10.

[352] Issue #7064: 3.3 epic: complete the post-acceptance commercial lifecycle. bokelley (AdCP lead architect), adcontextprotocol/adcp, 2026-08-30 (event 2026-09-21). https://github.com/adcontextprotocol/adcp/issues/7064 Accessed 2026-11-10.

[353] Issue #7068: Production evidence program: machine-observed AdCP implementation matrix. bokelley (AdCP lead architect), adcontextprotocol/adcp, 2026-08-30 (event 2026-11-08). https://github.com/adcontextprotocol/adcp/issues/7068 Accessed 2026-11-10.

[354] Roadmap (AdCP documentation, reference). AgenticAdvertising.org (AdCP), 2026-11-09. https://github.com/adcontextprotocol/adcp/blob/v3.2.9/docs/reference/roadmap.mdx Accessed 2026-11-10.

[355] Integral Ad Science to Provide Brand Safety and Suitability Measurement in ChatGPT Ads. Integral Ad Science (newsroom), 2026-10-05. https://integralads.com/news/ias-bss-chatgpt-ads/ Accessed 2026-11-10.

[356] DV Launches Attribution Measurement and Advances Brand Suitability for ChatGPT Ads. DoubleVerify (newsroom), 2026-10-05. https://doubleverify.com/company/newsroom/dv-launches-attribution-measurement-and-advances-brand-suitability-for-chatgpt-ads Accessed 2026-11-10.

[357] AdCP repository files at tag v3.3.0-beta.5 (docs/measurement/taxonomy.mdx; static/schemas/source/enums/performance-standard-metric.json; static/compliance/source/specialisms listing). AgenticAdvertising.org (AdCP) on GitHub, 2026-11-07. https://github.com/adcontextprotocol/adcp/tree/v3.3.0-beta.5 Accessed 2026-11-10.

[358] Meta launches brand suitability controls and third party verification for feed powered by AI (with dated updates). Meta (Meta for Business news), 2023-03-30 (event updates 2023-03-30 to 2024-10-24). https://www.facebook.com/business/news/brand-safety-suitability-feed-control-verification Accessed 2026-11-10.

[359] Brand suitability. Google (Display & Video 360 Help). https://support.google.com/displayvideo/answer/3032915?hl=en Accessed 2026-11-10.

[360] Activate pre-bid segments on Google Display & Video 360 | Adloox (Learning Center). Adloox, undated (footer © 2025). https://www.adloox.com/learning-center/google-dv360 Accessed 2026-11-10.

[361] Adloox site search index (Framer searchIndex JSON referenced by the live adloox.com pages), entries for legacy pages: /press, the DV360 guide of 23 November 2022 and the MFA DV360 guide of 15 April 2024. Adloox, undated (current site build). https://framerusercontent.com/sites/3xXsjLLbECSj0XYtEwmuu0/searchIndex-AwimyG20rifI.json Accessed 2026-11-10.

[362] BODACC open-data records for RCS Paris 513 450 544 (ADSAFE, then ADLOOX, Scope3 France, P39 France): 20 notices, 2009 to 2026. Bulletin officiel des annonces civiles et commerciales (DILA, French government), 2009-2026. https://bodacc-datadila.opendatasoft.com/api/explore/v2.1/catalog/datasets/annonces-commerciales/records?where=registre%20like%20%22513450544%22%20or%20registre%20like%20%22513%20450%20544%22&order_by=dateparution%20desc&limit=50 Accessed 2026-11-10.

[363] TAG Registry entry: Adloox. Trustworthy Accountability Group (TAG), undated live registry. https://hubgloo.opxar.com/company/52706971739 Accessed 2026-11-10.

[364] Adloox launches Brand Suitability & Safety Verification & Measurement for Meta, offering coverage for advertisers and agencies. Adloox, 2024-10-24. https://www.adloox.com/blog/adloox-launches-brand-suitability-safety-verification-measurement-for-meta-offering-coverage-for-advertisers-and-agencies Accessed 2026-11-10.

[365] Adloox Unveils New Segments & Unique Granularity in Google's Display & Video 360. Adloox, 2024-09-23. https://www.adloox.com/blog/adloox-unveils-new-segments-unique-granularity-in-google-display-video-360 Accessed 2026-11-10.

[366] Policy for Accreditation Representations by MRC-Accredited Services. Media Rating Council (MRC), 2024-02-15. https://mediaratingcouncil.org/sites/default/files/News/Process-Related/021524_Accreditation_Representation_Policy.pdf Accessed 2026-11-10.

[367] Made for Advertising sites (MFA): everything you need to know | Adloox (MFA center). Adloox, undated (latest dated item listed: 5 November 2024). https://www.adloox.com/platform/mfa-center Accessed 2026-11-10.

[368] Adalytics sitemap (blog URLs with lastmod stamps). Adalytics Research LLC. https://adalytics.io/sitemap.xml Accessed 2026-11-10.

[369] Adloox for publishers. Adloox, undated (footer © 2025). https://www.adloox.com/solutions/publishers Accessed 2026-11-10.

[370] Terms & Conditions | Adloox (website terms of use V 1.1, with 'How personal data is used for the Adloox SOLUTION'). Adloox (Adloox SAS), 2022-09-08. https://www.adloox.com/legal/terms-and-conditions Accessed 2026-11-10.

[371] Third-party measurement on YouTube (Display & Video 360 Help). Google. https://support.google.com/displayvideo/answer/13853322?hl=en Accessed 2026-11-10.

[372] Brand Protection Solutions | Protect Your Brand on TikTok. TikTok For Business, 2026 (undated). https://ads.tiktok.com/business/en/brand-safety-suitability Accessed 2026-11-10.

[373] Programmatic Partners. Peer39, undated (footer © 2026). https://www.peer39.com/programmatic-partners Accessed 2026-11-10.

[374] OpenPath Publisher Terms and Conditions v.10.31.24. The Trade Desk, Inc., 2024-11. https://assets.thetradedesk.com/assets/documents/legal/OpenPath-Publisher-Terms-and-Conditions-v.10.31.24.pdf Accessed 2026-11-10.

[375] Décision n° 23-MC-01 du 4 mai 2023 relative à une demande de mesures conservatoires de la société Adloox (version publique). Autorité de la concurrence (France), 2023-05-04. https://www.autoritedelaconcurrence.fr/sites/default/files/integral_texts/2023-05/23mc01.pdf Accessed 2026-11-10.

[376] What is the Open Measurement SDK and how it serves CTV advertising. Adloox, 2023-06-15. https://www.adloox.com/blog/open-measurement-sdk-for-ctv-advertising-iab-tech-lab Accessed 2026-11-10.

[377] AdCP documentation corpus (llms-full.txt). AgenticAdvertising.org (AdCP), undated (live corpus). https://docs.adcontextprotocol.org/llms-full.txt Accessed 2026-11-10.

[378] Peer39 Acquires Adloox from Scope3, Unifying Contextual Intelligence, Brand Safety Across Google, Meta, and the Open Web. Peer39 (distributed by EIN Presswire; media contact Sharp Pen Media), 2026-06-02. https://www.einpresswire.com/article/916779905/peer39-acquires-adloox-from-scope3-unifying-contextual-intelligence-brand-safety-across-google-meta-and-the-open-web Accessed 2026-11-10.

[379] Adloox expands its capabilities to Meta, rolling out Viewability Measurement for Facebook, Instagram and the Meta Audience Network. Adloox, 2024-02-20. https://www.adloox.com/blog/adloox-expands-its-capabilities-to-meta-rolling-out-viewability-measurement-for-facebook-instagram-and-the-meta-audience-network Accessed 2026-11-10.

[380] What DV360's Brand Suitability Change Means for Campaigns (blog). David Simutis, Peer39, 2026-10-30. https://www.peer39.com/blog/dv360-brand-safety-changes Accessed 2026-11-10.

[381] Adloox Limited (08701942): Annual Report and Unaudited Financial Statements for the Period from 1 July 2025 to 31 December 2025. Adloox Limited, filed at Companies House (UK), 2026-09-30. https://find-and-update.company-information.service.gov.uk/company/08701942/filing-history/MzU0ODU3ODc2N2FkaXF6a2N4/document?format=xhtml&download=1 Accessed 2026-11-10.

[382] YouTube Activation Partner Roster (Enterprise Marketing Portal). Google, undated. https://enterprisemarketingportal.google/find-youtube-activation-partners Accessed 2026-11-10.

[383] YouTube launches specialized activation partners program. PPC Land (Luis Rijo), 2025-10-06. https://ppc.land/youtube-launches-specialized-activation-partners-program/ Accessed 2026-11-10.

[384] TAG Registry entry: Channel Factory LLC. Trustworthy Accountability Group (TAG). https://hubgloo.opxar.com/company/52772278965 Accessed 2026-11-10.

[385] kidSAFE member page: Channel Factory's ViewIQ Solution. kidSAFE Seal Program. https://kidsafe.com/member/channelfactory_vendor Accessed 2026-11-10.

[386] Brand Safety (capabilities page). Channel Factory, 2026. https://channelfactory.com/capabilities/brand-safety Accessed 2026-11-10.

[387] Brand Suitability (Channel Factory capabilities page). Channel Factory. https://channelfactory.com/capabilities/brand-suitability Accessed 2026-11-10.

[388] Your Trusted Kid YouTube Partner Overview. Channel Factory, 2026. https://channelfactory.com/resources/your-trusted-kid-youtube-partner-overview Accessed 2026-11-10. *Access: partial extract.*

[389] Conscious Business Program. Channel Factory, 2026. https://channelfactory.com/cbp Accessed 2026-11-10.

[390] Channel Factory home page. Channel Factory, 2026. https://channelfactory.com/ Accessed 2026-11-10.

[391] Google Drive folder linked as the PwC report ('Download the Report'). Channel Factory (link target). https://drive.google.com/drive/u/0/folders/12P36_F9do5zeSmcwX1Zg_JoKSRy3Hog1 Accessed 2026-11-10. *Access: partial extract.*

[392] Channel Factory home page (archived copy). Channel Factory, 2023 (event 2023-06-09). https://web.archive.org/web/20230609222228/https://channelfactory.com/ Accessed 2026-11-10.

[393] YouTube Measurement Program (archived copy, 10 July 2024). YouTube (Google), undated (captured 2024-07-10) (event 2024-07-10). https://web.archive.org/web/20240710222350/https://www.youtube.com/measurementprogram/ Accessed 2026-11-10.

[394] About TikTok's Video Exclusion List (TikTok Business Help Center). TikTok, 2026-02. https://ads.tiktok.com/help/article/about-video-exclusion-list Accessed 2026-11-10.

[395] Channel Factory Launches AI Slop Detection. Channel Factory, 2026-10-02. https://channelfactory.com/resources/channel-factory-launches-ai-slop-detection Accessed 2026-11-10.

[396] Channel Factory Launches Premium CTV Intelligence Layer, Bringing Contextual Precision to the Living Room. Channel Factory (GlobeNewswire), 2026-09-28. https://www.globenewswire.com/news-release/2026/09/28/3369940/0/en/channel-factory-launches-premium-ctv-intelligence-layer-bringing-contextual-precision-to-the-living-room.html Accessed 2026-11-10.

[397] Premium CTV, Now With Context (blog). Channel Factory, 2026-10-15. https://channelfactory.com/resources/premium-ctv-now-with-context Accessed 2026-11-10.

[398] Additional policies for derived metrics and data storage (YouTube API Services). Google (YouTube for Developers), 2026-09-14. https://developers.google.com/youtube/terms/derived-metrics-policy Accessed 2026-11-10.

[399] Portfolio (Truelink Capital). Truelink Capital, 2026. https://truelinkcap.com/portfolio/ Accessed 2026-11-10.

[400] About TikTok's solution for invalid traffic with DoubleVerify (TikTok Business Help Center). TikTok, 2026-09. https://ads.tiktok.com/help/article/about-tiktok-doubleverify-invalid-traffic-solution Accessed 2026-11-10.

[401] Approved third-party providers (Amazon Ads policy). Amazon Ads, 2026-06-01 (last updated). https://advertising.amazon.com/resources/ad-policy/approved-3p-ad-servers Accessed 2026-11-10.

[402] DoubleVerify Legal Center: Hosted Main Service Agreement (v2.0, effective 11 February 2026), Partner-Pay Rates and Partner-Pay Solution Descriptions (v3.0, effective 21 May 2026), DV Publisher Suite SLAs (v1.0, effective 11 February 2026), EULA (v2.0, effective 1 May 2025). DoubleVerify (hosted on PactSafe), 2026-05-21 (latest document version). https://vault.pactsafe.io/s/20034d0c-1f43-4e41-b070-97502a449455/legal.html Accessed 2026-11-10.

[403] DoubleVerify Launches Content-Level Controls on Meta, Powering Media Quality and Advertising Effectiveness (release). DoubleVerify, 2025-02-18. https://doubleverify.com/company/newsroom/doubleverify-launches-content-level-controls-on-meta-powering-media-quality-and-advertising-effectiveness Accessed 2026-11-10.

[404] Media Rating Council Accreditations (DV MRC one-sheet, PDF). DoubleVerify, 2026-04-08. https://doubleverify.com/hubfs/content/MRC%20PDF/DV_MRC_One-Sheet.pdf Accessed 2026-11-10.

[405] DV Fraud Verification: Protect Your Brand Against Bad Actors (product page). DoubleVerify, 2026. https://doubleverify.com/products/advertisers/verify/fraud Accessed 2026-11-10.

[406] DV Brand Suitability: Protect Brand Equity (product page). DoubleVerify, 2026. https://doubleverify.com/products/advertisers/verify/brand-suitability Accessed 2026-11-10.

[407] DV Viewability: Make Sure Your Ad Can Be Seen (product page). DoubleVerify, 2026. https://doubleverify.com/products/advertisers/verify/viewability Accessed 2026-11-10.

[408] Addressing Publisher Measurement with Impartiality. DoubleVerify, 2024-06-13. https://doubleverify.com/blog/omnichannel/dvmap/addressing-publisher-measurement-with-impartiality Accessed 2026-11-10.

[409] DoubleVerify Impact Report 2023. DoubleVerify, 2023. https://doubleverify.com/hubfs/assets/pdfs/DV_ImpactReport_2023.pdf Accessed 2026-11-10.

[410] DoubleVerify Launches DV Authentic Streaming TV at CES 2026 (release). DoubleVerify, 2026-01-06. https://doubleverify.com/company/newsroom/doubleverify-launches-dv-authentic-streaming-tv-at-ces-2026 Accessed 2026-11-10.

[411] DoubleVerify Takes Down First Large-Scale Ad Impression Fraud Scheme in Audio (release). DoubleVerify, 2023-02-02. https://doubleverify.com/company/newsroom/doubleverify-takes-down-first-large-scale-ad-impression-fraud-scheme-in-audio Accessed 2026-11-10.

[412] DoubleVerify Launches AI-Powered Brand Suitability Reporting for YouTube Audio Ads Campaigns (release). DoubleVerify, 2026-06-11. https://doubleverify.com/company/newsroom/doubleverify-launches-ai-powered-brand-suitability-reporting-for-youtube-audio-ads-campaigns-expanding-transparency-in-listening-first-environments Accessed 2026-11-10.

[413] DoubleVerify Holdings, Inc. Form 8-K (Item 7.01, expected closing timing). DoubleVerify Holdings, Inc. (SEC filing), 2026-11-09. https://www.sec.gov/Archives/edgar/data/1819928/000110465926124574/tm2629362d1_8k.htm Accessed 2026-11-09.

[414] Letter from Senator Mark R. Warner to FTC Chairman Andrew Ferguson. U.S. Senator Mark R. Warner, 2025-03-24. https://www.warner.senate.gov/wp-content/uploads/media/doc/250324.Warner_Ferguson_Bot_fraud_letter.FINAL.SIGNED.pdf Accessed 2026-11-10.

[415] Digital Platforms Accreditation Status (chart of large digital platforms), updated 22 October 2026. Media Rating Council (MRC), 2026-10-22. https://mediaratingcouncil.org/sites/default/files/Digital%20Platforms%20Chart/Digital%20Platforms%20Accreditation%20Status_10.22.26.xlsx.pdf Accessed 2026-11-10.

[416] Accreditation letter to FreeWheel, Inc. (GI092024_FreeWheel Accreditation Letter): FreeWheel Analytics platform. Media Rating Council (George W. Ivie), 2024-09-20. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Freewheel/GI092024_FreeWheel%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[417] Media Rating Council accreditation updates, January 2024 through March 2024. Media Rating Council, 2024-05-01. https://mediaratingcouncil.org/sites/default/files/News/Accreditation-Matters/050124_MRC_Accreditation_Status_Communication.pdf Accessed 2026-11-10.

[418] Closing the Loop: MediaGuard + FraudSensor. HUMAN Security, 2026-10-29. https://docs.humansecurity.com/advertising/closing-the-loop-mediaguard-and-fraudsensor Accessed 2026-11-10.

[419] IVT Taxonomy. HUMAN Security, 2026-09-15. https://docs.humansecurity.com/advertising/ivt-taxonomy Accessed 2026-11-10.

[420] TAG Registry entry: HUMAN (registry widget behind tagtoday.net/registry). Trustworthy Accountability Group, 2026-11-10. https://hubgloo.opxar.com/search?q=HUMAN Accessed 2026-11-10.

[421] Brand Safety (HUMAN solution page), Internet Archive capture of 4 February 2026. HUMAN Security, Inc., 2026-02-04. https://web.archive.org/web/20260204234045id_/https://www.humansecurity.com/platform/solutions/brand-safety/ Accessed 2026-11-10.

[422] Ad Fraud (HUMAN solution page), Internet Archive capture of 9 February 2026. HUMAN Security, Inc., 2026-02-09. https://web.archive.org/web/20260209082151id_/https://www.humansecurity.com/platform/solutions/ad-fraud/ Accessed 2026-11-10.

[423] HUMAN Solution Brief: MRC Accreditation (PDF). HUMAN Security, Inc., 2023-08-16. https://www.humansecurity.com/wp-content/uploads/HUMAN_Solution-Brief_MRC-Accreditation-2.pdf Accessed 2026-11-10.

[424] HUMAN Service Terms and Conditions (updated May 29, 2025), Internet Archive capture of 28 March 2026. HUMAN Security, Inc., 2025-05-29. https://web.archive.org/web/20260328221144id_/https://www.humansecurity.com/service-terms-and-conditions/ Accessed 2026-11-10.

[425] HUMAN Subscription Agreement (Standard Terms; updated August 5, 2025), Internet Archive capture of 4 January 2026. HUMAN Security, Inc., 2025-08-05. https://web.archive.org/web/20260104083750id_/https://www.humansecurity.com/subscription-agreement/ Accessed 2026-11-10.

[426] HUMAN Authorized Third Party Data Use Agreement (updated January 29, 2025), Internet Archive capture of 14 September 2025. HUMAN Security, Inc., 2025-01-29. https://web.archive.org/web/20250914145045id_/https://www.humansecurity.com/authorized-third-party-data-use-agreement/ Accessed 2026-11-10.

[427] HUMAN documentation sitemap and index (docs.humansecurity.com/sitemap.xml; /advertising/llms.txt), with responses for linked pages. HUMAN Security, 2026-11-10. https://docs.humansecurity.com/sitemap.xml Accessed 2026-11-10.

[428] Letter from Senator Mark R. Warner to Attorney General Pamela Bondi. U.S. Senator Mark R. Warner, 2025-03-24. https://www.warner.senate.gov/public/_cache/files/2/a/2a67af93-4d82-4c4f-af37-cf20b77a0c7f/01C0BAE9B8B322E4E8A5AD3D45822A96E5E8CAC2E3513A20206BB082B2A5C6D2.250324.warner-bondi-bot-fraud-letter.final.signed.pdf Accessed 2026-11-10.

[429] Warner Pushes Federal Trade Commission, Justice Department to Address Rampant Fraud in Digital Advertising (press release). Office of U.S. Senator Mark R. Warner, 2025-03-28 (event 2025-03-24). https://www.warner.senate.gov/newsroom/press-releases/warner-pushes-federal-trade-commission-justice-department-to-address-rampant-fraud-in-digital-advertising/ Accessed 2026-11-10.

[430] Magnite Elevates Inventory Quality with HUMAN: Ad Fraud Defense Case Study (PDF). HUMAN Security, Inc., 2025-12-01. https://www.humansecurity.com/wp-content/uploads/2025/12/HUMAN_Case-Study_Magnite.pdf Accessed 2026-11-10.

[431] HUMAN website sitemaps (sitemap_index, page, post, newsroom and resource sitemaps). HUMAN Security, Inc., 2026-11-10. https://www.humansecurity.com/sitemap_index.xml Accessed 2026-11-10.

[432] HUMAN Security Disrupts CTV Device Spoofing Operation "NewsJunkie". HUMAN Security, Inc. (GlobeNewswire), 2026-07-07. https://www.globenewswire.com/news-release/2026/07/07/3323281/0/en/human-security-disrupts-ctv-device-spoofing-operation-newsjunkie.html Accessed 2026-11-09.

[433] About TikTok's solution for invalid traffic with Integral Ad Science (TikTok Business Help Center). TikTok, 2026-11. https://ads.tiktok.com/help/article/about-tiktok-integral-ad-science-invalid-traffic-solution Accessed 2026-11-10.

[434] IAS's Evolving Partnership with X (IAS Insider). Integral Ad Science, 2026-03-02. https://integralads.com/insider/ias-evolving-partnership-x/ Accessed 2026-11-10.

[435] TAG Registry search result: Integral Ad Science. Trustworthy Accountability Group, 2026-11-10. https://hubgloo.opxar.com/search?q=Integral%20Ad%20Science Accessed 2026-11-10.

[436] Inconsistencies in Classification of Online News Articles: A Call for Common Standards in Brand Safety Services (arXiv HTML v1). Smith, Grossman, Torres-Agüero, Sen, Borcea and Chen (NJIT; DeepSee.io), 2026-01-03. https://arxiv.org/html/2601.01303v1 Accessed 2026-11-10.

[437] Brand Safety & Suitability Solutions (solutions page). Integral Ad Science, 2026-06-10. https://integralads.com/solutions/brand-safety-suitability/ Accessed 2026-11-10.

[438] Mobile in-app solutions page. Integral Ad Science, 2026-06-10. https://integralads.com/solutions/mobile-in-app/ Accessed 2026-11-10.

[439] Ad Fraud Protection & Detection Solutions (product page). Integral Ad Science, 2026. https://integralads.com/solutions/ad-fraud/ Accessed 2026-11-10.

[440] Integral Ad Science, Inc. Standard Terms and Conditions. Integral Ad Science, 2026-07-13. https://integralads.com/terms-and-conditions/ Accessed 2026-11-10.

[441] IAS Help Center article titles and summaries (helpcenter.integralplatform.com: sitemap and 15 article pages). Integral Ad Science, 2026-11-10. https://helpcenter.integralplatform.com/sitemap.xml Accessed 2026-11-10. *Access: partial extract.*

[442] IAS Launches IAS Total TV Solutions Giving Marketers 'Linear-Like' Transparency for Connected TV (IAS Insider). Integral Ad Science, 2026-04-27. https://integralads.com/insider/ias-launches-total-tv-solutions-for-connected-tv/ Accessed 2026-11-10.

[443] IAS Launches First AI-Driven, Independent Brand Safety and Suitability Measurement for Meta Threads. Integral Ad Science, 2025-10-16. https://integralads.com/news/ias-launches-brand-safety-and-suitability-measurement-for-meta-threads/ Accessed 2026-11-10.

[444] TikTok Selects IAS for Brand Safety Measurement Across Pangle Advertisers (IAS Insider). Integral Ad Science, 2025-10-28. https://integralads.com/insider/tiktok-selects-ias-for-brand-safety-measurement-across-pangle-advertisers/ Accessed 2026-11-10.

[445] Expanding Transparency: IAS Brings Brand Safety and Suitability Measurement to YouTube Audio Ads Campaigns (IAS Insider). Integral Ad Science, 2026-06-11. https://integralads.com/insider/ias-brings-brand-safety-and-suitability-to-youtube-audio-ads/ Accessed 2026-11-10.

[446] IAS Brings Episode-Level Pre-Bid Optimization for Podcasts on and off Spotify via The Trade Desk (IAS Insider). Integral Ad Science, 2026-06-16. https://integralads.com/insider/ias-episode-level-pre-bid-optimization-for-podcasts/ Accessed 2026-11-10.

[447] Letter from Senators Blackburn and Blumenthal to Lisa Utzschneider (Integral Ad Science). U.S. Senators Marsha Blackburn and Richard Blumenthal, 2025-02-07. https://www.blackburn.senate.gov/services/files/0F377EF5-0A54-47D2-BB29-431DC5A5F0DE Accessed 2026-11-10.

[448] IAS Description of Services (Description of IAS Solutions; Additional Terms). Integral Ad Science, 2026-07-13. https://integralads.com/description-of-solutions/ Accessed 2026-11-10.

[449] Integral Ad Science (IAS) portfolio company page. Novacap, undated. https://novacapcorp.com/company/integral-data-science/ Accessed 2026-11-10.

[450] Static Media portfolio company page. Novacap, undated. https://novacapcorp.com/company/static-media/ Accessed 2026-11-10.

[451] Form 15 (15-12G): Certification and Notice of Termination of Registration, Integral Ad Science Holding Corp.. Integral Ad Science Holding Corp. (SEC filing), 2026-01-02. https://www.sec.gov/Archives/edgar/data/1842718/000114036126000141/ef20061704_15-12g.htm Accessed 2026-11-09.

[452] EDGAR company filing index, Integral Ad Science Holding Corp. (CIK 0001842718), 40 most recent filings. U.S. Securities and Exchange Commission (EDGAR), 2026-11-10. https://www.sec.gov/cgi-bin/browse-edgar?action=getcompany&CIK=1842718&type=&dateb=&owner=include&count=40 Accessed 2026-11-10.

[453] How MGID Detects and Prevents Invalid Traffic (IVT) (MGID Help). MGID, undated (footer 2026). https://help.mgid.com/how-mgid-detects-and-prevents-invalid-traffic-ivt Accessed 2026-11-10.

[454] Pixalate documentation: Amazon DSP. Pixalate, Inc., undated (live page). https://docs.pixalate.com/data-feeds/amazon-dsp Accessed 2026-11-10.

[455] Amazon DSP Enables Pixalate Pre-bid Ad Fraud Blocking. Pixalate, Inc., 2022-01-27. https://www.pixalate.com/blog/amazon-dsp-enables-pixalate-ad-fraud-blocking Accessed 2026-11-10.

[456] Pixalate documentation: IVT Taxonomy 2.0. Pixalate, Inc., undated (live page). https://docs.pixalate.com/trust/ivt-taxonomy Accessed 2026-11-10.

[457] Pixalate documentation: Analytics FAQ. Pixalate, Inc., undated (live page). https://docs.pixalate.com/analytics/faq Accessed 2026-11-10.

[458] Pixalate documentation: IP block lists (IPv4 and IPv6). Pixalate, Inc., undated (live page). https://docs.pixalate.com/data-feeds/ip-block-lists Accessed 2026-11-10.

[459] Pixalate Master Subscription Terms (effective February 12, 2026). Pixalate, Inc., 2026-02-12. https://www.pixalate.com/hubfs/Finalized_2026-02-12_Pixalate%20Master%20Subscription%20Terms%20(MST)%20(1).pdf Accessed 2026-11-10.

[460] Pixalate Agreements and Terms (legal page listing 13 Master Subscription Terms versions). Pixalate, Inc., undated (live page). https://www.pixalate.com/legal Accessed 2026-11-10.

[461] Pixalate documentation changelog (index and entries of 17 Nov 2025, 2 Dec 2025, 5 Feb 2026, 11 Mar 2026, 20 Aug 2026, 28 Aug 2026, 23 Sep 2026 and 14 Oct 2026). Pixalate, Inc., 2026-10-14 (event 2025-11-17 to 2026-10-14). https://docs.pixalate.com/changelog/llms.txt Accessed 2026-11-10.

[462] Pixalate documentation: Clawback report. Pixalate, Inc., undated (live page). https://docs.pixalate.com/analytics/clawback-report Accessed 2026-11-10.

[463] Pixalate introduces 'Clawback Report: Ad Fraud Dispute Resolution System for SSPs and DSPs'. Pixalate, Inc., 2025-11-04. https://www.pixalate.com/blog/pixalate-introduces-clawback-report-ad-fraud-dispute-resolution-system-for-ssps-and-dsps Accessed 2026-11-10.

[464] Pixalate documentation: MRT FAQ. Pixalate, Inc., undated (live page). https://docs.pixalate.com/mrt/faq Accessed 2026-11-10.

[465] Analytics (product page). Pixalate, Inc., undated (live page). https://www.pixalate.com/products/analytics/ Accessed 2026-11-10.

[466] Q1 2026 Ad Fraud Benchmarks Report for EMEA (release, with the global Q1 2026 IVT rates). Pixalate, 2026-04-29. https://www.pixalate.com/blog/q1-2026-ad-fraud-benchmarks-report-for-emea Accessed 2026-11-09.

[467] Pixalate documentation: Blocking FAQ. Pixalate, Inc., undated (live page). https://docs.pixalate.com/data-feeds/faq Accessed 2026-11-10.

[468] MRC-accredited measurement areas (Pixalate documentation). Pixalate. https://docs.pixalate.com/trust/mrc-measurement-areas Accessed 2026-11-10.

[469] Pixalate documentation: Video viewability and engagement (beta). Pixalate, Inc., undated (live page). https://docs.pixalate.com/analytics/video-viewability Accessed 2026-11-10.

[470] pixalate.com sitemap (2,470 URLs). Pixalate, Inc., 2026-11-10. https://www.pixalate.com/sitemap.xml Accessed 2026-11-10.

[471] Pixalate Docs index (llms.txt; 118 entries incl. API reference). Pixalate, Inc., 2026-11-10. https://docs.pixalate.com/llms.txt Accessed 2026-11-10.

[472] Pixalate's Q3 2026 Recap: Product Launches, Research, and Industry Highlights. Pixalate, Inc., 2026-10-02. https://www.pixalate.com/blog/pixalates-q3-2026-recap-product-launches-research-and-industry-highlights Accessed 2026-11-10.

[473] Pixalate Launches Ad Fraud API v3.0, Returns Ad Fraud Reason Codes for 46 IVT Types. Pixalate, Inc., 2026-10-13. https://www.pixalate.com/blog/pixalate-launches-ad-fraud-api-v3.0 Accessed 2026-11-10.

[474] Pixalate documentation: FAQ. Pixalate. https://docs.pixalate.com/get-started/faq Accessed 2026-11-10.

[475] Pixalate documentation: Seller Trust Index (STI 2.0). Pixalate, 2026-04-01. https://docs.pixalate.com/mrt/seller-trust-index Accessed 2026-11-10.

[476] Pixalate Unveils Ad Fraud 'Clawback Warnings' Across 17.5M+ Programmatic Publishers. Pixalate, Inc., 2026-10-06. https://www.pixalate.com/blog/pixalate-unveils-ad-fraud-clawback-warnings-across-17.5m-programmatic-publishers Accessed 2026-11-10.

[477] Pixalate documentation: Get support. Pixalate, Inc., undated (live page). https://docs.pixalate.com/get-started/support Accessed 2026-11-10.

[478] Vidillion, Inc. v. Pixalate, Inc., No. 2:18-cv-07270 (C.D. Cal.), docket. CourtListener (RECAP), 2018-08-17 (event 2018-08-17 to 2020-01-16). https://www.courtlistener.com/docket/7689046/vidillion-inc-v-pixalate-in

c/ Accessed 2026-11-10.

[479] Order Granting Motion for Judgment on the Pleadings (Dkt. No. 43), Vidillion v. Pixalate, ECF 46. U.S. District Court, C.D. Cal. (Judge Dale S. Fischer), 2019-03-15. https://storage.courtlistener.com/recap/gov.uscourts.cacd.720149/gov.uscourts.cacd.720149.46.0.pdf Accessed 2026-11-10.

[480] Order Granting Motion to Dismiss (Dkt. No. 54), Vidillion v. Pixalate, ECF 62. U.S. District Court, C.D. Cal. (Judge Dale S. Fischer), 2019-06-05. https://storage.courtlistener.com/recap/gov.uscourts.cacd.720149/gov.uscourts.cacd.720149.62.0.pdf Accessed 2026-11-10.

[481] Judgment, Vidillion v. Pixalate, ECF 77. U.S. District Court, C.D. Cal. (Judge Dale S. Fischer), 2019-08-07. https://storage.courtlistener.com/recap/gov.uscourts.cacd.720149/gov.uscourts.cacd.720149.77.0.pdf Accessed 2026-11-10.

[482] Pixalate documentation: Pre-bid data feeds. Pixalate, Inc., undated (live page). https://docs.pixalate.com/data-feeds/overview Accessed 2026-11-10.

[483] Criteo Partners with Pixalate to Provide Additional Protection for Advertisers' Investments. Pixalate, Inc., 2019-10-10. https://www.pixalate.com/blog/criteo-partnership-protect-advertising-investments Accessed 2026-11-10.

[484] Pixalate Launches 'Made For Advertising' (MFA) Detection and Blocking Technology for Websites, Connected TV & Mobile Apps. Pixalate, Inc., 2023-11-21. https://www.pixalate.com/blog/pixalate-launches-made-for-advertising-mfa-detection-and-blocking-technology-for-websites-mobile-ctv Accessed 2026-11-10.

[485] Advertising Week 2024: New Brand Suitability Controls (with updates of February 13, 2025 and March 27, 2025). Meta (Meta for Business news), 2024-10-10. https://www.facebook.com/business/news/Advertising-Week-2024-New-Brand-Suitability-Controls Accessed 2026-11-10.

[486] About TikTok's Video Exclusion List (TikTok Business Help Center). TikTok, 2026-02. https://ads.tiktok.com/help/article/about-video-exclusion-list?lang=en Accessed 2026-11-10.

[487] TikTok Brand Safety & Suitability Playbook: Navigate with Confidence. TikTok for Business, 2025-04-14. https://ads.tiktok.com/business/library/TikTokBrandSafetySuitabilityPlaybook.pdf Accessed 2026-11-10.

[488] Zefr acquires Israeli AI firm Adverif.ai, bolstering technology-led approach to identifying and defunding misinformation. Zefr (PR Newswire, reposted on zefr.com), 2022-07-20. https://zefr.com/press/zefr-acquires-israeli-ai-firm-adverif-ai-bolstering-technology-led-approach-to-identifying-and-defunding-misinformation Accessed 2026-11-09.

[489] Zefr Named to TikTok Marketing Partner Program. Zefr (press page), 2023-02-23. https://zefr.com/press/zefr-named-to-tiktok-marketing-partner-program Accessed 2026-11-10.

[490] Zefr Announces Brand Safety & Suitability Measurement Capability for its Misinformation Category on YouTube (September 2024). Zefr (press page), 2024-09-26. https://zefr.com/press/zefr-announces-brand-safety-suitability-measurement-capability-for-its-misinformation-category-on-youtube Accessed 2026-11-10.

[491] Bringing More Transparency and Confidence to Advertisers on Google Video Partners. Zefr (press page), 2026-01-29. https://zefr.com/press/bringing-more-transparency-and-confidence-to-advertisers-on-google-video-partners Accessed 2026-11-10.

[492] Zefr and Reddit Partner to Give Advertisers Deeper Insights and More Precise Suitability Controls. Zefr (press page), 2026-09-30. https://zefr.com/press/zefr-and-reddit-partner-to-give-advertisers-deeper-insights-and-more-precise-suitability-controls Accessed 2026-11-10.

[493] Company (about page). Zefr, undated. https://zefr.com/about/company Accessed 2026-11-10.

[494] Zefr and TikTok Advance Innovation in AI-Powered Safety and Suitability (Search, Smart+, TikTok Lite, GMV Max). Zefr (press page), 2026-04-21. https://zefr.com/press/zefr-and-tiktok-advance-innovation-in-ai-powered-safety-and-suitability-expanding-coverage-to-search-upgraded-campaign-creation-experience-for-smart-and-gmv-max Accessed 2026-11-10.

[495] The Classification Engine (technology page). Zefr, 2026-03-30. https://zefr.com/technology/zefr-classification Accessed 2026-11-10.

[496] Standards & Suitability Framework (technology page). Zefr, 2026-03-30. https://zefr.com/technology/standards-suitability-framework Accessed 2026-11-10.

[497] Terms of Use (website terms of service). ZEFR, Inc., 2026. https://zefr.com/terms-of-services Accessed 2026-11-10.

[498] Support page. Zefr, 2026. https://zefr.com/support Accessed 2026-11-10.

[499] Zefr Media Services (solution page). Zefr, 2026-03-25. https://zefr.com/solution/media-services Accessed 2026-11-10.

[500] Zefr Expands TikTok Product to Provide Advertisers With Suitability Exclusions, in Collaboration with TikTok's Inventory Filter (PR Newswire, September 2023). Zefr (press page), 2023-09-21. https://zefr.com/press/zefr-expands-tiktok-product-to-provide-advertisers-with-suitability-exclusions-in-collaboration-with-tiktoks-inventory-filter-pr-newswire-september-2023 Accessed 2026-11-10.

[501] About (company snapshot, principles, methodology). Zefr, 2026-03-24. https://zefr.com/about Accessed 2026-11-10.

[502] Form D notice of exempt offering, ZEFR, INC. (accession 0001474764-16-000001). ZEFR, INC. (filed with the SEC), 2016-03-08 (event 2016-03-01). https://www.sec.gov/Archives/edgar/data/1474764/000147476416000001/xslFormDX01/primary_doc.xml Accessed 2026-11-10.

[503] How Elon Musk Took Down Brand Safety Group GARM - What's Next for Advertisers Looking for Controls (Ad Age, August 9, 2024, by Garett Sloane and Jack Neff; reposted by Zefr). Zefr (press page), 2024-08-09. https://zefr.com/press/how-elon-musk-took-down-brand-safety-group-garm-whats-next-for-advertisers-looking-for-controls Accessed 2026-11-10.

[504] Zefr Partners with Snap to Launch Brand Safety and Suitability Measurement Solution (October 2024). Zefr (press page), 2024-10-29. https://zefr.com/press/zefr-partners-with-snap-to-launch-brand-safety-and-suitability-measurement-solution-october-2024 Accessed 2026-11-10.

[505] Observation log of this profile: zefr.com/.well-known/brand.json, AgenticAdvertising.org member API, Reddit-Zefr microsite, CourtListener entry page. No Fluff Advisory research log (this profile), 2026-11-10. https://zefr.com/.well-known/brand.json Accessed 2026-11-10.

[506] Brand safety (X Help Center, business and advertising). X Corp., 2026. https://help.x.com/en/business-and-advertising/brand-safety Accessed 2026-11-10.

[507] Zefr Announces Launch of Pre-Screen Brand Safety Solution for Google's Search Partner Network (SPN) (ExchangeWire, May 2025, reposted). Zefr (press page), 2025-05-08. https://zefr.com/press/zefr-announces-launch-of-pre-screen-brand-safety-solution-for-googles-search-partner-network-spn-3 Accessed 2026-11-10.

[508] About third-party pre-screen program for the Search Partner Network (Google Ads Help). Google, 2026. https://support.google.com/google-ads/answer/17104980?hl=en Accessed 2026-11-10.

[509] Customers | Adloox. Adloox, undated (footer © 2025). https://www.adloox.com/customers Accessed 2026-11-10.

[510] Businesses could lose $16.4 billion to online advert fraud in 2017. CNBC, 2017-03-15. https://www.cnbc.com/2017/03/15/businesses-could-lose-164-billion-to-online-advert-fraud-in-2017.html Accessed 2026-11-10.

[511] IVT 2.0 Compliance Status: Accredited Services (as of October 29th, 2021). Media Rating Council, 2021-10-29. https://mediaratingcouncil.org/sites/default/files/News/Accreditation-Matters/MRC%20IVT%202.0%20Status%20Report%20on%20Accredited%20Vendors.pdf Accessed 2026-11-10.

[512] Adloox Integrations Directory | Meta. Adloox, undated (footer © 2025). https://www.adloox.com/integrations/meta Accessed 2026-11-10.

[513] The MRC Just Redefined Brand Safety: What It Means for Advertisers (blog). Dannalyn Prado, Peer39, 2025-10-28. https://www.peer39.com/blog/the-mrc-just-redefined-brand-safety-what-it-means-for-advertisers Accessed 2026-11-10.

[514] Adloox Tag Request. Peer39, undated (footer © 2026). https://www.peer39.com/adloox-tag-request Accessed 2026-11-10.

[515] How to Use 3rd Party Tags (Taboola developer docs, Pixel). Taboola (Realize), 2026-02-03. https://developers.taboola.com/pixel/docs/how-to-use-3rd-party-tags Accessed 2026-11-10.

[516] About us | Adloox. Adloox (Peer39), 2026. https://www.adloox.com/company Accessed 2026-11-10.

[517] Mobile & In-app | Adloox. Adloox, undated (footer © 2025). https://www.adloox.com/channels/mobile Accessed 2026-11-10.

[518] Programmatic | Adloox. Adloox, undated (footer © 2025). https://www.adloox.com/channels/programmatic Accessed 2026-11-10.

[519] Recherche d'entreprises API record for SIREN 513450544 (P39 FRANCE), with INSEE Sirene and RNE fields. Direction interministérielle du numérique (French government), data from INSEE Sirene and the INPI national company register, 2026-11-09. https://recherche-entreprises.api.gouv.fr/search?q=p39%20france&page=1&per_page=10 Accessed 2026-11-10.

[520] 5 questions to Brian Sykora, SVP, Product Strategy and Corporate Development at Jun Group. Adloox, 2024-05-07. https://www.adloox.com/blog/5-questions-to-brian-sykora-svp-product-strategy-and-corporate-development-at-jun-group Accessed 2026-11-10.

[521] Terms of Service (Updated: June 15th, 2020). Peer39 (P39 Tech, LLC), 2020-06-15. https://www.peer39.com/service-terms Accessed 2026-11-10.

[522] Privacy Policy. Peer39, undated. https://www.peer39.com/privacy-policy/ Accessed 2026-11-10.

[523] CourtListener RECAP party-name searches: 'Adloox' (2 results), 'Peer39' (0), 'P39' (1). Free Law Project (CourtListener), 2026. https://www.courtlistener.com/?type=r&party_name=Adloox Accessed 2026-11-10.

[524] Adloox integrates with Equativ. Adloox, 2023-01-25. https://www.adloox.com/blog/adloox-integrates-with-equativ Accessed 2026-11-10.

[525] Brut. partners with Adloox to strengthen advertising activities optimization. Adloox, 2023-09-13. https://www.adloox.com/blog/brut-partners-with-adloox-to-strengthen-advertising-activities-optimization Accessed 2026-11-10.

[526] Advanced Contextual Data for Publishers (Publisher Solutions). Peer39, undated (footer © 2026). https://www.peer39.com/publisher-solutions Accessed 2026-11-10.

[527] Scope3 acquires Adloox to bolster media sustainability solutions and insights across more media platforms. Scope3, 2024-11-05. https://sustainability.scope3.com/news/scope3-acquires-adloox-to-bolster-media-sustainability-solutions-and-insights-across-more-media-platforms Accessed 2026-11-10.

[528] Contextual Targeting (capabilities page). Channel Factory, 2026. https://channelfactory.com/capabilities/contextual-targeting Accessed 2026-11-10.

[529] New Study by MAGNA & Channel Factory Finds 'Misaligned' Content Erodes the Impact of Hardworking Ads. MAGNA (IPG Mediabrands), MAGNA Media Trials, in partnership with Channel Factory, 2021-09-28 (event 2021-09-22 (dateline)). https://magnaglobal.com/new-study-by-magna-channel-factory-finds-misaligned-content-erodes-the-impact-of-hardworking-ads/ Accessed 2026-11-09.

[530] The Proximity Effect: Quantifying the impact of misaligned content in the wild west of video. MAGNA Media Trials with Channel Factory, 2021-09. https://ipg-wp-media-mgl-glb.s3.us-east-2.amazonaws.com/magna/wp-content/uploads/2021/09/27162727/MAGNA-Channel-Factory-The-Proximity-Effect.pdf Accessed 2026-11-10.

[531] OMG - Annalect Research (resource page). Channel Factory, 2025-12-08. https://channelfactory.com/resources/omg-annalect-research Accessed 2026-11-10. *Access: partial extract.*

[532] About Us. Channel Factory, 2026. https://channelfactory.com/about-us Accessed 2026-11-10.

[533] Channel Factory Chosen as Launch Member of the YouTube Activation Partners program. Channel Factory (GlobeNewswire), 2025-10-08 (event 2025-10-06). https://www.globenewswire.com/news-release/2025/10/08/3163797/0/en/channel-factory-chosen-as-launch-member-of-the-youtube-activation-partners-program.html Accessed 2026-11-10.

[534] Free Content Strategy Analysis. Channel Factory, 2026. https://channelfactory.com/content-strategy-analysis Accessed 2026-11-10.

[535] AERO: Agentic Execution for ROI (landing page). Channel Factory, 2026. https://lp.channelfactory.com/aero Accessed 2026-11-10. *Access: partial extract.*

[536] Advertising Technology Partners Directory (in beta). Google, undated. https://enterprisemarketingportal.google/find-ad-tech-partners Accessed 2026-11-10.

[537] McDonald's Case Study. Channel Factory, 2026-09-11. https://channelfactory.com/resources/mcdonalds-case-study Accessed 2026-11-10. *Access: partial extract.*

[538] Privacy and Cookie Policy. Channel Factory, LLC, 2026-08-17. https://channelfactory.com/privacy-policy Accessed 2026-11-10.

[539] Form D, Notice of Exempt Offering of Securities: Channel Factory, LLC (accession 0001140361-18-031927). Channel Factory, LLC (SEC EDGAR), 2018-07-09 (event 2018-06-21). https://www.sec.gov/Archives/edgar/data/1745403/000114036118031927/xslFormDX01/primary_doc.xml Accessed 2026-11-10.

[540] Star Mountain Announces Exit of Channel Factory Investment. Star Mountain Capital, 2025-06-05. https://www.starmountain.com/media-center/channel-factory Accessed 2026-11-10.

[541] Media Platforms. Channel Factory, 2026. https://channelfactory.com/media-platforms Accessed 2026-11-10.

[542] DV Solutions for CTV (channel page). DoubleVerify, undated (live 2026-11-10). https://doubleverify.com/channels/ctv Accessed 2026-11-10.

[543] Did Verification Catch It? Exploring Made for Advertising Sites and DV's Measurement Standards (Transparency Center blog). DoubleVerify, 2024-06-27. https://doubleverify.com/blog/web/verify/did-verification-catch-it-exploring-made-for-advertising-sites-and-dvs-measurement-standards Accessed 2026-11-10.

[544] Unproven Brand Suitability Tools Threaten Advertiser Confidence in News. DoubleVerify (blog, News Accelerator), 2026-03-05. https://doubleverify.com/blog/web/verify/unproven-brand-suitability-tools-threaten-advertiser-confidence-in-news Accessed 2026-11-09.

[545] DoubleVerify Bolsters Media Authentication on Meta: Extends Brand Safety and Suitability to Facebook and Instagram Feeds & Reels (release). DoubleVerify, 2024-01-17. https://doubleverify.com/company/newsroom/doubleverify-bolsters-media-authentication-on-meta-extends-brand-safety-and-suitability-to-facebook-and-instagram-feeds-reels Accessed 2026-11-10.

[546] DoubleVerify Debuts Industry-leading In-Feed Brand Safety and Suitability Solution on Twitter (release). DoubleVerify, 2023-01-25. https://doubleverify.com/company/newsroom/doubleverify-debuts-industry-leading-in-feed-brand-safety-and-suitability-solution-on-twitter Accessed 2026-11-10.

[547] DoubleVerify Expands Brand Safety and Suitability Measurement to YouTube Shorts (release). DoubleVerify, 2023-12-12. https://doubleverify.com/company/newsroom/doubleverify-expands-brand-safety-and-suitability-measurement-to-youtube-shorts Accessed 2026-11-10.

[548] DoubleVerify Expands Its AI-Powered Brand Safety & Suitability Offering for TikTok to Include Pre-Bid Video Controls (release). DoubleVerify, 2025-04-14. https://doubleverify.com/company/newsroom/doubleverify-expands-its-ai-powered-brand-safety-suitability-offering-for-tiktok-to-include-pre-bid-video-controls Accessed 2026-11-10.

[549] Approved Vendors and Supported Products (Snapchat Business Help Center). Snap Inc.. https://businesshelp.snapchat.com/s/article/third-party-tracking-vendors?language=en_US Accessed 2026-11-10.

[550] Set up DoubleVerify measurement for CTV Ads (LinkedIn Marketing Solutions Help). LinkedIn (Microsoft). https://www.linkedin.com/help/lms/answer/a10273720 Accessed 2026-11-10.

[551] DV's Code of Ethics. DoubleVerify. https://doubleverify.com/code-of-ethics Accessed 2026-11-10.

[552] DoubleVerify Expands Access to Open Web URL-Level Transparency (blog). DoubleVerify, 2025-02-25. https://doubleverify.com/blog/web/verify/doubleverify-expands-access-to-open-web-url-level-transparency Accessed 2026-11-10.

[553] DoubleVerify's Response to The Wall Street Journal's Fact Check and Adalytics Report. DoubleVerify, 2025-03-28. https://doubleverify.com/blog/web/verify/doubleverifys-response-to-the-wall-street-journals-fact-check-and-adalytics-report Accessed 2026-11-10.

[554] Metrics (Media Security). HUMAN Security, 2026-08-20. https://docs.humansecurity.com/advertising/metrics-media-security Accessed 2026-11-10.

[555] Brand Safety Web Crawler (GlancelotBot). HUMAN Security, 2026-06-29. https://docs.humansecurity.com/advertising/brand-safety-web-crawler Accessed 2026-11-10.

[556] Open Web Community Guidelines (made-for-advertising indicators). HUMAN Security, 2026-10-23. https://docs.humansecurity.com/advertising/community_site_guidelines Accessed 2026-11-10.

[557] MediaGuard API: Perform a lookup request (v4). HUMAN Security, 2026-10-27. https://docs.humansecurity.com/advertising/reference/perform-lookup-request-4 Accessed 2026-11-10.

[558] FraudSensor Domain Mismatch IVT Classifications (and FAQ). HUMAN Security, 2026-06-29. https://docs.humansecurity.com/advertising/fs-domain-mismatch-ivt-classifications Accessed 2026-11-10.

[559] Viewability (HUMAN solution page), Internet Archive capture of 21 February 2026. HUMAN Security, Inc., 2026-02-21. https://web.archive.org/web/20260221061204id_/https://www.humansecurity.com/platform/solutions/viewability/ Accessed 2026-11-10.

[560] HUMAN Security Solutions to Enhance Protections for Ad Traffic Quality on LinkedIn (press release dated June 18, 2024), Internet Archive capture of 4 February 2026. HUMAN Security, Inc., 2024-06-18. https://web.archive.org/web/20260204173514id_/https://www.humansecurity.com/newsroom/linkedin-harnesses-the-power-of-human-security-to-protect-feed-from-bad-bots-and-ad-fraud/ Accessed 2026-11-10.

[561] United States v. Google LLC, 1:23-cv-00108 (E.D. Va.), Document 1132-2 (Exhibit B), pages 141 to 143 of 325: "2017 Strategy Paper: 3rd party Fraud Verification Service". Google LLC (document produced in litigation; filed 2024-08-06), 2024-08-06 (event 2017). https://storage.courtlistener.com/recap/gov.uscourts.vaed.533508/gov.uscourts.vaed.533508.1132.2_1.pdf Accessed 2026-11-10.

[562] The Trade Desk Partners with White Ops to Become First Advertising Platform to Block Fraudulent Impressions Before They Are Purchased (press release), Internet Archive capture of 13 October 2024. White Ops and The Trade Desk (GlobeNewswire text on HUMAN's newsroom), 2017-08-31. https://web.archive.org/web/20241013002914id_/https://www.humansecurity.com/newsroom/the-trade-desk-partners-with-white-ops-to-become-first-advertising-platform-to-block-fraudulent-impressions-before-they-are-purchased Accessed 2026-11-10.

[563] HUMAN Raises $50+ Million in Growth Funding to Protect the Digital Customer Journey and Defend Against Bots, Fraud and Risk. HUMAN Security, 2024-10-09. https://web.archive.org/web/20241107162102id_/https://www.humansecurity.com/newsroom/platform-expansion-fuel-the-future-of-human Accessed 2026-11-09.

[564] HUMAN Raises $100 Million In Growth Funding Round Led By WestCap and NightDragon. HUMAN Security, 2022-01-04. https://web.archive.org/web/20221208144311id_/https://www.humansecurity.com/newsroom/human-raises-100-million-in-growth-funding-round-led-by-westcap-and-nightdragon Accessed 2026-11-09.

[565] White Ops Announces Acquisition by Goldman Sachs Merchant Banking, ClearSky Security, and NightDragon. White Ops (now HUMAN Security), 2020-12-23. https://web.archive.org/web/20210331085037id_/https://www.humansecurity.com/newsroom/white-ops-announces-acquisition-by-goldman-sachs-merchant-banking-clearsky-security-and-nightdragon Accessed 2026-11-09.

[566] Form D: Singularity TopCo LLC (accession 0001944306-24-000003). Singularity TopCo LLC (filer) / SEC EDGAR, 2024-10-09 (event 2024-09-26). https://www.sec.gov/Archives/edgar/data/1944306/000194430624000003/primary_doc.xml Accessed 2026-11-10.

[567] HUMAN Ad Viewability Measurement Earns MRC Accreditation, Redefining Transparency, Explainability, and Trust in Ad Measurement. HUMAN Security, Inc. (release distributed by GlobeNewswire), 2026-04-27 (event 2026-04-14 (MRC letter)). https://www.globenewswire.com/news-release/2026/04/27/3281661/0/en/human-ad-viewability-measurement-earns-mrc-accreditation-redefining-transparency-explainability-and-trust-in-ad-measurement.html Accessed 2026-11-10.

[568] Form D: Singularity TopCo LLC (accession 0001944306-24-000002). Singularity TopCo LLC (filer) / SEC EDGAR, 2024-05-07 (event 2024-04-22). https://www.sec.gov/Archives/edgar/data/1944306/000194430624000002/primary_doc.xml Accessed 2026-11-10.

4000002/primary_doc.xml Accessed 2026-11-10.

[569] EDGAR company search: "human security", "singularity topco", "singularity buyer", "white ops". SEC EDGAR, 2026-11-10. https://www.sec.gov/cgi-bin/browse-edgar?company=singularity+topco&owner=include&count=40&action=getcompany Accessed 2026-11-10.

[570] Viewability solutions page. Integral Ad Science, 2025-04-16. https://integralads.com/solutions/viewability/ Accessed 2026-11-10.

[571] Media Quality Report: 21st Edition (global). Integral Ad Science, 2026-07. https://iabhongkong.com/sites/default/files/2026-07/GLB_MQR_21st_Edition_IAS.pdf Accessed 2026-11-09.

[572] Site indexing policy (IAS Crawler and IAS Wombles). Integral Ad Science, 2023-10-26. https://integralads.com/ias-privacy-data-management/policies/site-indexing-policy/ Accessed 2026-11-10.

[573] Why Integral Ad Science page. Integral Ad Science, 2026-08-12. https://integralads.com/why-ias/ Accessed 2026-11-10.

[574] Digital audio solutions page. Integral Ad Science, 2025-04-24. https://integralads.com/solutions/digital-audio/ Accessed 2026-11-10.

[575] Introducing the Industry's First Video-Level Brand Suitability Measurement for CTV (IAS Insider). Integral Ad Science, 2024-01-18. https://integralads.com/insider/video-level-brand-suitability-measurement-ctv/ Accessed 2026-11-10.

[576] IAS's Approach to Brand Safety & Suitability (IAS Insider). Integral Ad Science, 2025-03-24. https://integralads.com/insider/approach-to-brand-safety-suitability/ Accessed 2026-11-10.

[577] IAS Announces First-to-Market Meta Optimization Solution for Advertisers (Content Block Lists). Integral Ad Science, 2024-10-10. https://integralads.com/news/ias-meta-optimization-solution/ Accessed 2026-11-10.

[578] Our next step in Marketplace Quality: SIVT credits (blog post), Internet Archive capture of 6 June 2023. The Trade Desk, 2018-05-01. https://web.archive.org/web/20230606060203id_/https://www.thetradedesk.com/us/our-next-step-in-marketplace-quality-sivt-credits Accessed 2026-11-10.

[579] Integral Ad Science and Seedooh Partner to Bring Trusted Digital Metrics to OOH (release). Integral Ad Science, 2018-03-29. https://integralads.com/news/integral-ad-science-seedooh-partner-bring-trusted-digital-metrics-ooh/ Accessed 2026-11-10.

[580] Connected TV (CTV) Measurement Solutions (CTV & Video solutions page). Integral Ad Science. https://integralads.com/solutions/ctv-video/ Accessed 2026-11-10.

[581] MRC Accredited Services: Digital, Internet Archive capture of 16 June 2026. Media Rating Council (via Internet Archive), 2026-06-16. https://web.archive.org/web/20260616055022id_/https://mediaratingcouncil.org/accreditation/digital Accessed 2026-11-10.

[582] Samsung Ads and Publica by IAS Announce Multi-Year Extension of Exclusive Global Partnership to Power Global CTV Ad Serving (release). Integral Ad Science, 2025-09-30. https://integralads.com/news/samsung-ads-and-publica-announce-extension-2025/ Accessed 2026-11-10.

[583] IAS Quality Connect Brings Transparency to Publisher Campaign Optimization (IAS Insider). Integral Ad Science, 2026-06-18. https://integralads.com/insider/ias-quality-connect-publisher-campaign-transparency/ Accessed 2026-11-10.

[584] Cadent portfolio company page. Novacap, undated. https://novacapcorp.com/company/cadent/ Accessed 2026-11-10.

[585] IAS Signal Platform Media Partner End User License Agreement. Integral Ad Science, 2025-05-09. https://integralads.com/ias-legal-portal/ias-signal-platform-media-partner-end-user-license-agreement/ Accessed 2026-11-10.

[586] IAS Service Level Requirements. Integral Ad Science, 2024-01-10. https://integralads.com/ias-service-level-requirements/ Accessed 2026-11-10.

[587] Detect and Avoid Made For Advertising Sites with IAS's AI-Driven MFA Technology (IAS Insider). Integral Ad Science, 2024-03-05. https://integralads.com/insider/detect-avoid-made-for-advertising-sites-ai-driven-tec

hnology/ Accessed 2026-11-10.

[588] Mastercard Improves Campaign Performance with IAS. Integral Ad Science (case study), 2023-10-20. https://integralads.com/insider/mastercard-dentsu-case-study/ Accessed 2026-11-09. *Access: partial extract.*

[589] Pixalate documentation: CTV and OTT reporting. Pixalate, Inc., undated (live page). https://docs.pixalate.com/analytics/ctv-and-ott Accessed 2026-11-10.

[590] Pixalate documentation: Fraud API. Pixalate, Inc., undated (live page). https://docs.pixalate.com/apis/fraud-api Accessed 2026-11-10.

[591] Pixalate documentation: High Risk Domains. Pixalate, Inc., undated (live page). https://docs.pixalate.com/data-feeds/high-risk-domains Accessed 2026-11-10.

[592] Pixalate documentation: Metrics. Pixalate, Inc., undated (live page). https://docs.pixalate.com/analytics/metrics Accessed 2026-11-10.

[593] Pixalate documentation: MRC-accredited metrics. Pixalate, Inc., undated (live page). https://docs.pixalate.com/trust/mrc-accredited-metrics Accessed 2026-11-10.

[594] Pixalate documentation: MFA methodology and definitions. Pixalate, Inc., undated (live page). https://docs.pixalate.com/compliance/mfa-methodology Accessed 2026-11-10.

[595] Pixalate documentation: Glossary (Media Ratings Terminal). Pixalate, Inc., undated (live page). https://docs.pixalate.com/mrt/glossary Accessed 2026-11-10.

[596] 'Monarch': An apparent Connected TV/OTT ad fraud scheme stealing from premium brands and political ad budgets. Pixalate, 2020-03-25. https://www.pixalate.com/blog/monarch-ott-ctv-ad-fraud-scheme Accessed 2026-11-10.

[597] Pixalate documentation: COPPA FAQ. Pixalate, Inc., undated (live page). https://docs.pixalate.com/compliance/coppa-faq Accessed 2026-11-10.

[598] Pixalate documentation: Choose an integration method. Pixalate, Inc., undated (live page). https://docs.pixalate.com/analytics/integration-methods Accessed 2026-11-10.

[599] Pixalate documentation: Website tag blocking (Google Ad Manager). Pixalate, Inc., undated (live page). https://docs.pixalate.com/data-feeds/website-tag-blocking Accessed 2026-11-10.

[600] Global Invalid Traffic (IVT) Benchmarks Report, Q2 2026 (landing page). Pixalate, 2026 (undated page; footer reads Copyright 2025). https://www.pixalate.com/invalid-traffic-benchmarks-report-global Accessed 2026-11-09. *Access: partial extract.*

[601] Pixalate, Inc. Form D, Notice of Exempt Offering of Securities (accession 0001582620-13-000001). Pixalate, Inc. (SEC EDGAR, CIK 0001582620), 2013-07-29 (event 2013-06-11 (first sale)). https://www.sec.gov/Archives/edgar/data/1582620/000158262013000001/xslFormDX01/primary_doc.xml Accessed 2026-11-10.

[602] Pixalate becomes first company to receive MRC accreditation for OTT invalid traffic detection and filtration. Pixalate, Inc., 2018-11-14. https://www.pixalate.com/blog/mrc-accreditation-ott-invalid-traffic-detection-filtration Accessed 2026-11-10.

[603] Verve reaffirms commitment to fraud-free media with Pixalate partnership and TAG recertification. Pixalate, Inc., 2019-05-14. https://www.pixalate.com/blog/verve-partnership-tag-recertification-programmatic-ad-fraud Accessed 2026-11-10.

[604] Pixalate documentation: Privacy, brand safety and COPPA reports. Pixalate, Inc., undated (live page). https://docs.pixalate.com/analytics/compliance-reports Accessed 2026-11-10.

[605] TAG Registry entry: Pixalate. Trustworthy Accountability Group (registry served from hubgloo.opxar.com), undated (live entry). https://hubgloo.opxar.com/company/52808694440 Accessed 2026-11-10.

[606] Pixalate documentation: Welcome to Pixalate. Pixalate, Inc., undated (live page). https://docs.pixalate.com/get-started/welcome Accessed 2026-11-10.

[607] Pixalate Terms of Use (effective 24 November 2025). Pixalate Europe Limited and its affiliates, 2025-11-24. https://www.pixalate.com/terms Accessed 2026-11-10.

[608] Pixalate Announces $18.1 Million Additional Growth Capital to Accelerate Global Expansion and Media Ratings Platform. Pixalate, Inc., 2021-09-02. https://www.pixalate.com/blog/funding-announcement-2021 Accessed 2026-11-10.

[609] Pre-bid Fraud Protection with Blocklists (Blocking product page). Pixalate, Inc., undated (live page). https://www.pixalate.com/products/blocking/ Accessed 2026-11-10.

[610] Pixalate documentation: Allowlisting policy. Pixalate, Inc., undated (live page). https://docs.pixalate.com/trust/allowlisting Accessed 2026-11-10.

[611] Pixalate documentation: MRC accreditation. Pixalate, Inc., undated (live page). https://docs.pixalate.com/trust/mrc-accreditation Accessed 2026-11-10.

[612] Viewability (solution page). Zefr, undated. https://zefr.com/solution/viewability Accessed 2026-11-10.

[613] Zefr home page. Zefr, undated. https://zefr.com/ Accessed 2026-11-10.

[614] Zefr Expands Brand Suitability Protection on TikTok with AI-Powered Exclusion Lists (PR Newswire, April 2025). Zefr (press page), 2025-04-14. https://zefr.com/press/zefr-expands-brand-suitability-protection-on-tiktok-with-ai-powered-exclusion-lists-pr Accessed 2026-11-10.

[615] Zefr to Expand AI-Driven Brand Suitability Capabilities for Greater Control on Meta (PR Newswire, October 2024). Zefr (press page), 2024-10-10. https://zefr.com/press/zefr-to-expand-ai-driven-brand-suitability-capabilities-for-greater-control-on-meta Accessed 2026-11-10.

[616] Zefr Selected as Brand Safety Reporting Partner by YouTube. Zefr (press page), 2020-10-21. https://zefr.com/press/zefr-selected-as-brand-safety-reporting-partner-by-youtube Accessed 2026-11-10.

[617] Zefr Awarded New U.S. Patent for AI-Powered Content Classification Process. Zefr (press page), 2025-12-18. https://zefr.com/press/zefr-awarded-new-u.s.-patent-for-ai-powered-content-classification-process Accessed 2026-11-10.

[618] AI vs. Human Moderators: A Comparative Evaluation of Multimodal LLMs in Content Moderation for Brand Safety. Adi Levi, Or Levi, Sardhendu Mishra, Jonathan Morra (all Zefr Inc.), 2025-08-07. https://arxiv.org/abs/2508.05527 Accessed 2026-11-09.

[619] Zefr & Forrester Consulting Release Study on the Current State of Video Advertising and Brand Suitability ('The Maturation of Brand Suitability'). 4A's (research listing), 2022 (page reads 'Published 01.10.2022', day and month order ambiguous). https://www.aaaa.org/research-report/zefr-forrester-consulting-release-study-on-the-current-state-of-video-advertising-and-brand-suitability/ Accessed 2026-11-09.

[620] TikTok partners with Zefr on brand safety and brand suitability measurement (TikTok For Business post, January 24, 2022, reposted by Zefr). Zefr (press page), 2022-01-24. https://zefr.com/press/tiktok-partners-with-zefr-on-brand-safety-and-brand-suitability-measurement Accessed 2026-11-10.

[621] Zefr & Meta Expand AI-Powered Brand Suitability Measurement to the Instagram Feed (ExchangeWire PressBox, 2 August 2023, reposted). Zefr (press page), 2023-08-04. https://zefr.com/press/zefr-meta-expand-ai-powered-brand-suitability-measurement-to-the-instagram-feed Accessed 2026-11-10.

[622] Controls Pre vs Post Impressions (technology page). Zefr, 2026. https://zefr.com/technology/controls-pre-vs-post-impressions Accessed 2026-11-10.

[623] More Control, More Confidence: Zefr Launches Content Block Lists for Threads Feed. Zefr (press page), 2026-05-18. https://zefr.com/press/more-control-more-confidence-zefr-launches-content-block-lists-for-threads-feed-2 Accessed 2026-11-10.

[624] Zefr and YouTube Expand AI-Powered Suitability Insights to Audio Ads Campaigns. Zefr (press page), 2026-06-11. https://zefr.com/press/zefr-and-youtube-expand-ai-powered-suitability-insights-to-audio-ads-campaigns Accessed 2026-11-10.

[625] As Ad Dollars Flow into CTV, Brand Safety Considerations Follow: Zefr and Iris.TV partnering (Adweek article by Andrew Blustein, reposted by Zefr). Zefr (press page), 2020-11-20. https://zefr.com/press/zefr-and-iris-tv-are-partnering-to-bring-a-more-nuanced-approach-to-blocking-or-allowing-streaming-content Accessed 2026-11-10.

[626] Mazda case study (YouTube, CX-50 launch, 2021 campaign). Zefr, 2022-07-13. https://zefr.com/resources/case-studies/mazda Accessed 2026-11-10.

[627] CMI Media Group case study: Major Pharmaceutical Brand Leveraged Zefr Suitability Data on Meta. Zefr, 2023-09-14. https://zefr.com/resources/case-studies/cmi-media-group Accessed 2026-11-10.

[628] TAG Registry entry for ZEFR (registry widget served from hubgloo.opxar.com). Trustworthy Accountability Group (TAG), 2026. https://hubgloo.opxar.com/search?q=Zefr Accessed 2026-11-10.

[629] As Global Video Investment Accelerates, Zefr Continues Growth With Exclusive Mitsui Partnership In Japan (PR Newswire, April 28, 2021). Zefr (press page), 2021-04-28. https://zefr.com/press/zefr-mitsui-partnership-garm Accessed 2026-11-10.

[630] Zefr Launches Opti to Help Demand Gen and Performance Max Advertisers Optimize Media More Effectively. Zefr (press page), 2026-08-06. https://zefr.com/press/zefr-opti Accessed 2026-11-10.

[631] zefr.com sitemaps (index and ten section sitemaps, 291 URLs including 222 press pages). Zefr, 2026. https://zefr.com/sitemaps-2-sitemap.xml Accessed 2026-11-10.

[632] GARM Brand Safety Floor + Suitability Framework (knowledge page). WFA, 2022-06-17. https://wfanet.org/knowledge/item/2022/06/17/GARM-Brand-Safety-Floor--Suitability-Framework-3 Accessed 2026-11-10. *Access: partial extract.*

[633] Third-Party Accreditation, Certification and Review. Comscore, 2026. https://www.comscore.com/About/Third-Party-Review Accessed 2026-11-10.

[634] Zefr is the First Third-Party Partner to receive MRC Accreditation for Content-Level Brand Safety and Suitability Reporting on YouTube. Zefr, 2026-03-09. https://zefr.com/press/zefr-is-the-first-third-party-partner-to-receive-mrc-accreditation-for-content-level-brand-safety-and-suitability-reporting-on-youtube Accessed 2026-11-10.

[635] Company | Protected by Mediaocean. Protected Media (Mediaocean), 2026. https://www.protected.media/company Accessed 2026-11-10.

[636] MRC Accreditation Areas. Pixalate, 2025. https://www.pixalate.com/mrc-accreditation-areas Accessed 2026-11-10.

[637] Interactive Advertising Bureau Guidelines for the Conduct of Ad Verification. IAB and MRC, 2012-02-14. https://mediaratingcouncil.org/sites/default/files/Standards/Ad-Verification-Guideline-for-the-Conduct-of-2.pdf Accessed 2026-11-10.

[638] MRC Guidance on Existing Standards for AI Use in Media Measurement. Media Rating Council (MRC), 2026-07-08. https://mediaratingcouncil.org/sites/default/files/Standards/MRC_Existing_Guidance_on_AI_Use_in_Media_Measurement_FINAL.pdf Accessed 2026-11-10.

[639] Other Guidance. Media Rating Council (MRC), 2026. https://mediaratingcouncil.org/standards-and-guidelines/information Accessed 2026-11-10.

[640] MRC accreditation decisions report: Q3-24 (July-September 2024). Media Rating Council (MRC), 2024-10-11. https://mediaratingcouncil.org/sites/default/files/News/MRC%20Accreditation%20Status%20Quarterly%20Update%20Q3%2024_101124.pdf Accessed 2026-11-10.

[641] MRC accreditation decisions report: Q4-24 (October-December 2024). Media Rating Council (MRC), 2025-02-13. https://mediaratingcouncil.org/sites/default/files/News/MRC%20Accreditation%20Status%20Quarterly%20Update_Q4%2024_2.12_25%20.pdf Accessed 2026-11-10.

[642] DV Earns MRC Accreditation for CTV Viewability, Reinforcing its Leadership in Pre- and Post-Bid CTV Measurement. DoubleVerify, 2024-04-24. https://doubleverify.com/company/newsroom/dv-earns-mrc-accreditation-for-ctv-viewability-reinforcing-its-leadership-in-pre-and-post-bid-ctv-measurement Accessed 2026-11-10.

[643] Quality Impressions (product page). Integral Ad Science, 2026. https://integralads.com/innovation/quality-impressions/ Accessed 2026-11-10.

[644] Certified Against Fraud Guidelines. TAG Anti-Fraud Working Group, 2026-07. https://www.tagtoday.net/hubfs/CAF/TAG%20CAF%20Guidelines%20Final.pdf Accessed 2026-11-10.

[645] Changelog for Open Measurement SDK JavaScript clients (CHANGELOG.md). IAB Tech Lab (GitHub: InteractiveAdvertisingBureau/Open-Measurement-JSClients), 2026-11-02. https://github.com/InteractiveAdvertisingBureau/Open-Measurement-JSClients/blob/master/CHANGELOG.md Accessed 2026-11-10.

[646] ads.cert. IAB Tech Lab, 2024-02-14. https://iabtechlab.com/standards/ads-cert/ Accessed 2026-11-10.

[647] ads.cert (standards page). IAB Tech Lab, 2024-02-14. https://iabtechlab.com/ads-cert/ Accessed 2026-11-10.

[648] Public Release of Ads.cert: Securing the CTV Supply Chain (blog). Rob Hazan, Index Exchange, on the IAB Tech Lab blog, 2022-01-27. https://iabtechlab.com/blog/public-release-ads-cert-securing-ctv-supply-chain/ Accessed 2026-11-10.

[649] IABTechLab/adscert repository metadata and commit history (GitHub API). IAB Tech Lab (GitHub) (event 2023-03-08 (last commit)). https://api.github.com/repos/IABTechLab/adscert/commits?per_page=3 Accessed 2026-11-10.

[650] Certified Against Malvertising Guidelines. TAG, 2026-07. https://www.tagtoday.net/hubfs/CAM/TAG%20CAM%20Guidelines.pdf Accessed 2026-11-10.

[651] Accredited Services: Digital, Internet Archive captures from 2024-02-23 to 2026-10-06. Internet Archive (of MRC page), 2024-02-23 to 2026-10-06. https://web.archive.org/web/20250906131329/https://mediaratingcouncil.org/accreditation/digital Accessed 2026-11-10.

[652] Ad Fraud Prevention & Monitoring Platform (home page). Escalated Solutions, LLC, 2026. https://escalated.io/ Accessed 2026-11-10. *Access: partial extract.*

[653] Escalated.io Releases S8FTY, A Patent Pending Real-Time Invalid Traffic Filtration System For Brands and Publishers. Escalated Solutions, LLC, 2026-04-01. https://escalated.io/n/Escalated.io-Releases-S8FTY-A-Patent-Pending-Real-Time-Invalid-Traffic-Filtration-System-For-Brands-and-Publishers.html Accessed 2026-11-10.

[654] TAG Certification Programs. Trustworthy Accountability Group (TAG), 2026. https://www.tagtoday.net/certifications Accessed 2026-11-10.

[655] Leading Trade Groups Define "Made For Advertising" Websites (press release). ANA, with the 4A's, WFA and ISBA, 2023-09-26. https://www.ana.net/content/show/id/81828 Accessed 2026-11-10.

[656] IAS's 20th Media Quality Report reveals Ad Fraud is on the Rise, But Optimization Strategies Protect Budgets and Deliver Higher Performance. Integral Ad Science, 2025-05-20. https://integralads.com/news/media-quality-report-20th-edition/ Accessed 2026-11-09.

[657] IAS's 21st Media Quality Report Shows How Proactive Optimization Turns Media Quality into a Performance Engine. Integral Ad Science, 2026-07-09. https://integralads.com/news/ias-21st-media-quality-report/ Accessed 2026-11-09.

[658] US Ad Industry Extends Success in Fight Against Fraud with Fourth Straight Year of Sub-1% IVT in TAG-Certified Channels. TAG (Trustworthy Accountability Group); study conducted by The 614 Group, 2024-10-29. https://www.tagtoday.net/pressreleases/us-ad-industry-extends-success-in-fight-against-fraud-with-fourth-straight-year-of-sub-1-ivt-in-tag-certified-channels Accessed 2026-11-09.

[659] TAG anti-fraud programme page (Certified Against Fraud; fraud research). TAG, undated (read 2026-11-09). https://www.tagtoday.net/fraud Accessed 2026-11-09.

[660] Cross-Industry Anti-Fraud Efforts Saved Advertisers $10.8 Billion in 2023. TAG (PR Newswire), 2024-10-09. https://prnewswire.com/news-releases/cross-industry-anti-fraud-efforts-saved-advertisers-10-8-billion-in-2023--302271031.html Accessed 2026-11-09.

[661] Humans make better content cops than AI, but cost 40x more (The Register, August 2025, as reposted by Zefr). Zefr press page reposting The Register, 2025-08. https://zefr.com/press/humans-make-better-content-cops-than-ai-but-cost-40x-more Accessed 2026-11-09.

[662] Brand Safety Spotlight: Leveraging Oracle Advertising Solutions to effectively monitor brand safety and improve performance (A Forrester Total Economic Impact Spotlight commissioned by Oracle). Forrester Consulting, commissioned by Oracle, 2022 (undated; cites TEI studies of November 2021 and February

2022). https://www.oracle.com/a/ocom/docs/forrester-brand-safety-spotlight-oracle-advertising.pdf Accessed 2026-11-09.

[663] Internet Archive CDX listing for adloox.com/blog/use-adloox-segment-on-google-display-video-360. Internet Archive, 2026. https://web.archive.org/cdx/search/cdx?url=adloox.com/blog/use-adloox-segment-on-google-display-video-360&output=txt&limit=20 Accessed 2026-11-10. *Access: partial extract.*

[664] DoubleVerify to Protect Advertisers Against 'Made for Advertising' Content with Industry First Holistic Solution (release). DoubleVerify, 2023-09-14 (event 2023-09-20). https://doubleverify.com/company/newsroom/doubleverify-to-protect-advertisers-against-made-for-advertising-content-with-industry-first-holistic-solution Accessed 2026-11-10.

[665] About Adalytics. Adalytics Research LLC. https://adalytics.io/about Accessed 2026-11-10.

[666] Letter from Senators Blackburn and Blumenthal to Mark Zagorski (DoubleVerify). U.S. Senators Marsha Blackburn and Richard Blumenthal, 2025-02-07. https://www.blackburn.senate.gov/services/files/F66FDFEE-380D-4C81-802A-2FD93626F8DB Accessed 2026-11-10.

[667] ANA Provides "First Look" at In-depth Programmatic Media Transparency Study (press release). ANA, 2023-06-19. https://www.ana.net/content/show/id/pr-2023-06-programmaticstudy Accessed 2026-11-10.

[668] ANA Releases Part 2 of Comprehensive Programmatic Media Buying Study (press release). ANA, 2023-12-05. https://www.ana.net/content/show/id/83522 Accessed 2026-11-10.

[669] Malvertising Defense FAQ. HUMAN Security (documentation site), undated. https://docs.humansecurity.com/advertising/malvertising-defense-faq Accessed 2026-11-09.

[670] Our next step in Marketplace Quality: SIVT credits (blog post, Internet Archive capture). The Trade Desk, Inc., 2018-05-01. https://web.archive.org/web/20230606060203/https://www.thetradedesk.com/us/our-next-step-in-marketplace-quality-sivt-credits Accessed 2026-11-10.

[671] Google LLC Advertising Program Terms (United States). Google LLC, 2025-07-29. https://payments.google.com/payments/apis-secure/u/0/get_legal_document?ldi=351439 Accessed 2026-11-10.

[672] Important updates to the Google Ads Terms of Service. Google (Google Ads Help). https://support.google.com/google-ads/answer/16875158?hl=en Accessed 2026-11-10.

[673] AdTrader, Inc. v. Google LLC, Order Granting Motion for Final Approval of Settlement and Granting Motion for Attorney Fees, Expenses, and Service Award (ECF No. 407). United States District Court, Northern District of California (Judge Beth Labson Freeman), 2022-11-01. https://cases.justia.com/federal/district-courts/california/candce/5:2017cv07082/320384/407/0.pdf Accessed 2026-11-10.

[674] Media Services Framework Agreement, RM6123 (REDACTED v2), with Framework Schedules and Lot 1 and Lot 2 Scope and Specification. Crown Commercial Service (now Government Commercial Agency), UK Government, 2021-12-14. https://assets.crowncommercial.gov.uk/wp-content/uploads/RM6123-Media-Services-framework-agreement-REDACTED-v2-1.odt Accessed 2026-11-10.

[675] Media Services (RM6123) agreement page. Government Commercial Agency (formerly Crown Commercial Service). https://www.gca.gov.uk/agreements/RM6123 Accessed 2026-11-10.

[676] WFA releases full version of Global Media Charter to reform digital ad ecosystem. World Federation of Advertisers (WFA), 2018-06-21. https://wfanet.org/knowledge/item/2018/06/21/WFA-releases-full-version-of-Global-Media-Charter-to-reform-digital-ad-ecosystem Accessed 2026-11-10.

[677] Global Media Charter (page for the Global Media Charter 2023). WFA. https://wfanet.org/globalmediacharter Accessed 2026-11-10. *Access: partial extract.*

[678] TAG Launches Ad Industry's First Global Brand Safety Certification Program. TAG, 2020-09-10. https://www.tagtoday.net/pressreleases/brand-safety-certification-launches Accessed 2026-11-10.

[679] TAG Recognizes Leading Companies Across Ad Industry Worldwide for Achieving 2026 Certifications. TAG, 2026-03-05 (event page header March 4, 2026; dateline March 5, 2026). https://www.tagtoday.net/pressreleases/raising-the-bar-tag-applauds-2026-certified-leaders-elevating-industry-standards-for-fraud-brand-safety-malvertising-and-transparency Accessed 2026-11-10.

[680] FTC Alters Final Consent Order in Response to Public Comments, Preventing Coordination in Global Advertising Merger. US Federal Trade Commission, 2025-09-26. https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-alters-final-consent-order-response-public-comments-preventing-coordination-global-advertising Accessed 2026-11-10.

[681] FTC Prevents Anticompetitive Coordination in Global Advertising Merger. US Federal Trade Commission, 2025-06-23. https://www.ftc.gov/news-events/news/press-releases/2025/06/ftc-prevents-anticompetitive-coordination-global-advertising-merger Accessed 2026-11-10.

[682] Omnicom Group Inc. Form 10-K for the fiscal year ended 31 December 2025. Omnicom Group Inc., 2026-02-20. https://www.sec.gov/Archives/edgar/data/29989/000002998926000006/omc-20251231.htm Accessed 2026-11-10.

[683] WPP plc Form 20-F for the fiscal year ended 31 December 2025. WPP plc, 2026-03-19. https://www.sec.gov/Archives/edgar/data/806968/000162828026019643/wpp-20251231.htm Accessed 2026-11-10.

[684] Media Services Framework (2021). ISBA, 2021-07-13. https://www.isba.org.uk/knowledge/media-services-framework-2021 Accessed 2026-11-10. *Access: partial extract.*

[685] ANA Introduces Brand Safety and Suitability Metrics (Q3 2025 Programmatic Transparency Benchmark release). Association of National Advertisers (ANA), 2025-11-05. https://www.ana.net/content/show/id/pr-2025-11-transparency Accessed 2026-11-10.

[686] Rumble Inc v. World Federation of Advertisers, 7:24-cv-00115 (N.D. Tex.), docket. Free Law Project (CourtListener; PACER/RECAP), 2024-08-06 (event terminated 2025-08-13). https://www.courtlistener.com/docket/69018499/rumble-inc-v-world-federation-of-advertisers/ Accessed 2026-11-10. *Access: partial extract.*

[687] CourtListener RECAP search: case name contains DoubleVerify (6 results). Free Law Project (CourtListener). https://www.courtlistener.com/?type=r&q=caseName%3A(DoubleVerify)&order_by=dateFiled%20desc Accessed 2026-11-10.

[688] CourtListener RECAP search: case name contains "Integral Ad Science" (4 results). Free Law Project (CourtListener). https://www.courtlistener.com/?type=r&q=caseName%3A(%22Integral%20Ad%20Science%22)&order_by=dateFiled%20desc Accessed 2026-11-10.

[689] PubMatic announces fraud-free program (press release). PubMatic, Inc., 2017-11-07. https://pubmatic.com/news/pubmatic-announces-fraud-free-program/ Accessed 2026-11-10.

[690] New Brand Suitability Controls (Meta update of February 2025, as posted on Zefr's press page). Meta (text), posted by Zefr, 2025-02-13. https://zefr.com/press/advertising-week-2024-new-brand-suitability-controls-2 Accessed 2026-11-09.

[691] Business Third Party Partner Viewability Requests (Marketing API reference). Meta for Developers. https://developers.facebook.com/documentation/ads-commerce/marketing-api/reference/business/third_party_partner_viewability_requests Accessed 2026-11-10.

[692] DoubleVerify to Introduce Pre-screen Content Controls on Meta. DoubleVerify, 2024-10-10. https://doubleverify.com/company/newsroom/doubleverify-to-introduce-pre-screen-content-controls-on-meta-strengthening-brand-safety-suitability-and-media-performance Accessed 2026-11-10.

[693] Privacy checks in Ads Data Hub. Google for Developers. https://developers.google.com/ads-data-hub/guides/privacy-checks Accessed 2026-11-10.

[694] About TikTok Ads Manager supported third-party measurement (TikTok Business Help Center). TikTok, 2026-01. https://ads.tiktok.com/help/article/tiktok-ads-manager-supported-3rd-party-measurement Accessed 2026-11-10.

[695] DoubleVerify First Measurement Provider to Earn MRC Accreditation for TikTok Video Viewability Reporting. DoubleVerify, 2026-04-23. https://doubleverify.com/company/newsroom/doubleverify-first-measurement-provider-to-earn-mrc-accreditation-for-tiktok-video-viewability-reporting Accessed 2026-11-10.

[696] About TikTok's Profile Feed Exclusion List (TikTok Business Help Center). TikTok, 2025-04. https://ads.tiktok.com/help/article/about-profile-feed-exclusion-list Accessed 2026-11-10.

[697] Expanded End-to-End Coverage Across TikTok (IAS Insider). Integral Ad Science, 2024-10-08. https://integralads.com/insider/enhanced-brand-safety-suitability-tiktok/ Accessed 2026-11-10.

[698] Accreditation letter to Amazon Ads: first-party measurement of Amazon DSP traffic. Media Rating Council (George W. Ivie), 2026-05-29. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Amazon/GI052926_Amazon%20DSP%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[699] Third Party Tracking FAQs (Snapchat Business Help Center). Snap Inc.. https://businesshelp.snapchat.com/s/article/third-party-tracking-faqs?language=en_US Accessed 2026-11-10.

[700] Accreditation letter to Snap: Ads Manager platform. Media Rating Council (George W. Ivie), 2026-07-29. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Snap/GI072926_Snap%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[701] DoubleVerify Strengthens Media Quality Authentication on Snap with Brand Safety and Suitability Measurement. DoubleVerify, 2024-11-05. https://doubleverify.com/company/newsroom/doubleverify-strengthens-media-quality-authentication-on-snap-with-brand-safety-and-suitability-measurement Accessed 2026-11-10.

[702] Accreditation letter to Pinterest: Ads Manager platform. Media Rating Council (George W. Ivie), 2026-01-20. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Pinterest/GI012026%20Pinterest%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[703] Third-party tracking for LinkedIn Ads (LinkedIn Help). LinkedIn (Microsoft). https://www.linkedin.com/help/linkedin/answer/a423714 Accessed 2026-11-10.

[704] The Mechanics of Independent Social Media Measurement (DV Transparency Center blog). DoubleVerify, 2024-05-30. https://doubleverify.com/blog/social/verify/the-mechanics-of-independent-social-media-measurement Accessed 2026-11-10.

[705] About Brand Safety Hub (TikTok Business Help Center). TikTok, 2026-06. https://ads.tiktok.com/help/article/about-brand-safety-hub Accessed 2026-11-10.

[706] About TikTok invalid traffic solutions (TikTok Business Help Center). TikTok, 2025-04. https://ads.tiktok.com/help/article/about-tiktok-invalid-traffic-solutions Accessed 2026-11-10.

[707] Flashtalking by Mediaocean fully integrates verification to inform next-generation ad serving and optimization. Mediaocean, 2022-05-12. https://www.mediaocean.com/flashtalking-by-mediaocean-fully-integrates-verification-ad-serving Accessed 2026-11-09.

[708] Mediaocean to acquire Flashtalking, adding complementary solutions to power $200 billion in annualized media spend. Mediaocean, 2021-07-13. https://www.mediaocean.com/flashtalking-acquisition Accessed 2026-11-09.

[709] Mediaocean Completes Acquisition of Innovid - Creating the Premier Global Independent, Omnichannel Ad Tech Platform (Exhibit 99.1 to Innovid's 8-K of 13 February 2025). Mediaocean (release filed by Innovid), 2025-02-13. https://www.sec.gov/Archives/edgar/data/1835378/000114036125004373/ef20043552_ex99-1.htm Accessed 2026-11-10.

[710] Innovid Corp. Definitive Proxy Statement (Schedule 14A, merger with Mediaocean). Innovid Corp. (SEC filing), 2025-01-10 (event 2025-02-11). https://www.sec.gov/Archives/edgar/data/1835378/000114036125000684/ny20040318x2_defm14a.htm Accessed 2026-11-09. *Access: partial extract.*

[711] Cross-Platform Intelligence (technology page). Zefr, 2026. https://zefr.com/technology/cross-platform-intelligence-youtube-meta-tiktok-snapchat-google-ads-dv360 Accessed 2026-11-10.

[712] Décision 23-MC-01 du 04 mai 2023 (decision page, with 'Recours et Pourvois'). Autorité de la concurrence (France), 2023-05-04. https://www.autoritedelaconcurrence.fr/fr/decision/relative-une-demande-de-mesures-conservatoires-de-la-societe-adloox Accessed 2026-11-10.

[713] Accreditation letter to Meta Platforms: Facebook and Instagram Ads. Media Rating Council (George W. Ivie), 2026-10-02. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Meta/GI100226_Meta%20FB%20%26%20IG%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[714] Brand Suitability vs. Safety: Why Verification Isn't Always Enough. David Simutis, Peer39 (blog), 2026-08-11 (event updated 2026-08-18). https://www.peer39.com/blog/brand-suitability-vs-brand-safety Accessed 2026-11-09.

[715] DoubleVerify Launches Industry-First Streaming TV Products to Drive Greater Transparency and Elevate CTV Advertising Quality. DoubleVerify (Business Wire), 2025-11-06. https://www.businesswire.com/news/home/20251106767494/en/DoubleVerify-Launches-Industry-First-Streaming-TV-Products-to-Drive-Greater-Transparency-and-Elevate-CTV-Advertising-Quality Accessed 2026-11-10.

[716] OM SDK CTV Release 1.4: What is new?. IAB Tech Lab, 2022-08. https://iabtechlab.com/wp-content/uploads/2022/08/OMSDK_What-is-new-in-1.4_master.pdf Accessed 2026-11-10.

[717] OM SDK CTV Release 1.5: What is new?. IAB Tech Lab, 2024. https://iabtechlab.com/wp-content/uploads/2024/08/OMSDK_What-is-new-in-1.5_master.docx.pdf Accessed 2026-11-10.

[718] Changelog for Open Measurement SDK JavaScript clients. IAB Tech Lab (GitHub: InteractiveAdvertisingBureau/Open-Measurement-JSClients), 2026-11-02. https://raw.githubusercontent.com/InteractiveAdvertisingBureau/Open-Measurement-JSClients/master/CHANGELOG.md Accessed 2026-11-10.

[719] Open Measurement Enters CTV (webinar deck). IAB Tech Lab (presenters from Oracle Advertising, DoubleVerify, Google and IAB Tech Lab), 2022-08-03. https://iab.com/wp-content/uploads/2022/08/OMSDK-Enters-CTV.pdf Accessed 2026-11-10.

[720] VAST 4.3 (Video Ad Serving Template). IAB Tech Lab (GitHub: InteractiveAdvertisingBureau/VAST4.x), 2022-12. https://raw.githubusercontent.com/InteractiveAdvertisingBureau/VAST4.x/main/4.3.md Accessed 2026-11-10.

[721] Giving Fraudsters the Cold Shoulder: Inside the Largest Connected TV Bot Attack (ICEBUCKET). White Ops Satori Threat Intelligence and Research Team, 2020-04 (event peak January 2020). https://web.archive.org/web/20200416153729id_/https://www.whiteops.com/blog/giving-fraudsters-the-cold-shoulder-inside-the-largest-connected-tv-bot-attack Accessed 2026-11-10.

[722] ICEBUCKET: Inside the Largest CTV Ad Fraud Bot Operation (briefing deck). TAG (Trustworthy Accountability Group) and White Ops, 2020-04-30. https://www.tagtoday.net/hubfs/TAG%20WO%20Briefing%20re%20ICEBUCKET%20-%20FINAL%202020.4.30.pdf Accessed 2026-11-10.

[723] Oracle Exposes Largest CTV Ad Fraud Operation Ever (StreamScam) (PR Newswire release via Nasdaq). Oracle, 2020-12-17 (event 2020). https://www.nasdaq.com/press-release/oracle-exposes-largest-ctv-ad-fraud-operation-ever-2020-12-17 Accessed 2026-11-10.

[724] DoubleVerify Uncovers and Stops 'OctoBot' - Unprecedented CTV-Focused Fraud Scheme With Multiple Tentacles. DoubleVerify (GlobeNewswire), 2021-04-20 (event since November 2019). https://www.globenewswire.com/news-release/2021/04/20/2213475/0/en/DoubleVerify-Uncovers-and-Stops-OctoBot-Unprecedented-CTV-Focused-Fraud-Scheme-With-Multiple-Tentacles.html Accessed 2026-11-10.

[725] DoubleVerify Neutralizes 'SmokeScreen,' a New Global CTV Fraud Scheme Using Screensavers to Hijack Streaming Devices. DoubleVerify (Business Wire), 2021-08-12. https://www.businesswire.com/news/home/20210812005190/en/DoubleVerify-Neutralizes-SmokeScreen-a-New-Global-CTV-Fraud-Scheme-Using-Screensavers-to-Hijack-Streaming-Devices Accessed 2026-11-10.

[726] DoubleVerify Fraud Lab Report: CTV Fraudsters Impersonate Smart Refrigerators to Falsify CTV Traffic (LeoTerra). DoubleVerify (Business Wire), 2022-08-24 (event LeoTerra first identified July 2020). https://www.businesswire.com/news/home/20220824005237/en/DoubleVerify-Fraud-Lab-Report-CTV-Fraudsters-Impersonate-Smart-Refrigerators-to-Falsify-CTV-Traffic Accessed 2026-11-10.

[727] HUMAN Disrupts Digital Supply Chain Threat Actor Scheme Originating from China (BADBOX and PEACHPIT). HUMAN, 2023-10-04 (event ad-fraud module found November 2022). https://web.archive.org/web/20231130013811id_/https://www.humansecurity.com/newsroom/human-disrupts-digital-supply-chain-threat-actor-scheme-originating-from-china Accessed 2026-11-10.

[728] HUMAN Exposes BADBOX 2.0 Scheme Infecting 1 Million Off-Brand Android Open Source Project Devices. HUMAN, 2025-03-05. https://web.archive.org/web/20250318001544id_/https://www.humansecurity.com/newsroom/human-exposes-badbox-2-0-scheme/ Accessed 2026-11-10.

[729] Imitation is the sincerest form of fraudery (CopyCatz). White Ops Satori Threat Intelligence and Research Team, 2021-01. https://web.archive.org/web/20210113131126id_/https://www.whiteops.com/blog/imitation-is-the-sincerest-form-of-fraudery Accessed 2026-11-10.

[730] HUMAN Orchestrates Unprecedented Private Takedown, VASTFLUX. HUMAN, 2023-01-19. https://web.archive.org/web/20230119120637id_/https://www.humansecurity.com/newsroom/human-orchestrates-unprecedented-private-takedown-vastflux Accessed 2026-11-10.

[731] HUMAN Discovers Konfety Ad Fraud Operation Wielding Novel 'Evil Twin' Evasion Method. HUMAN, 2024-07-16. https://web.archive.org/web/20240717095454id_/https://www.humansecurity.com/newsroom/human-discovers-konfety-ad-fraud-operation-wielding-novel-evil-twin-evasion-method Accessed 2026-11-10.

[732] Home Internet Connected Devices Facilitate Criminal Activity (Public Service Announcement I-060525-PSA). Federal Bureau of Investigation, Internet Crime Complaint Center (IC3), 2025-06-05. https://www.ic3.gov/PSA/2025/PSA250605 Accessed 2026-11-10.

[733] Google LLC v. Does 1-25, Complaint for Damages and Injunctive Relief (Document 22). Google LLC (plaintiff), U.S. District Court for the Southern District of New York, 1:25-cv-04503-JPO, 2025-07-11 (event complaint first filed under seal; TRO 2025-05-30). https://storage.courtlistener.com/recap/gov.uscourts.nysd.643466/gov.uscourts.nysd.643466.22.0.pdf Accessed 2026-11-10.

[734] Google LLC v. Does 1-25, Preliminary Injunction Order (Document 18). U.S. District Court for the Southern District of New York (1:25-cv-04503-JPO), 2025-06-27 (event TRO 2025-05-30, extended 2025-06-06). https://cases.justia.com/federal/district-courts/new-york/nysdce/1%3A2025cv04503/643466/18/0.pdf Accessed 2026-11-10.

[735] DoubleVerify and Roku Bring Greater Trust and Performance to Streaming TV. DoubleVerify and Roku, 2025-11-05. https://doubleverify.com/company/newsroom/doubleverify-and-roku-bring-greater-trust-and-performance-to-streaming-tv Accessed 2026-11-10.

[736] Why CTV App Publishers should ensure they are providing Open Measurement with OM SDK (info sheet). IAB Tech Lab, 2025-11. https://iabtechlab.com/wp-content/uploads/2025/11/OM-SDK-from-IAB-Tech-Lab-Info-Sheet-Nov-2025.pdf Accessed 2026-11-10.

[737] Australian Competition Tribunal: Current matters. Australian Competition Tribunal, 2026-11-10. https://www.competitiontribunal.gov.au/current-matters Accessed 2026-11-10.

[738] DoubleVerify Reports Third Quarter 2026 Financial Results (Exhibit 99.1 to 8-K of 30 October 2026). DoubleVerify Holdings, Inc., 2026-10-30. https://www.sec.gov/Archives/edgar/data/1819928/000110465926120670/dv-ex99d1.htm Accessed 2026-11-09.

[739] Form 25-NSE: Notification of Removal from Listing, Integral Ad Science Holding Corp. common stock. Nasdaq Stock Market LLC (SEC filing), 2025-12-23. https://www.sec.gov/Archives/edgar/data/1842718/000135445725001293/xslF25X02/primary_doc.xml Accessed 2026-11-09.

[740] EDGAR company filing list, Integral Ad Science Holding Corp. (CIK 0001842718). U.S. Securities and Exchange Commission (EDGAR), 2026-11-10. https://www.sec.gov/cgi-bin/browse-edgar?action=getcompany&CIK=1842718&type=&dateb=&owner=include&count=10 Accessed 2026-11-10.

[741] Innovid Corp. Form 8-K (Item 1.01, Agreement and Plan of Merger with Mediaocean LLC). Innovid Corp. (SEC filing), 2024-11-21. https://www.sec.gov/Archives/edgar/data/1835378/000121390024100750/ea0222073-8k_innovid.htm Accessed 2026-11-09.

[742] Innovid Corp. Form 8-K (completion of the merger with Mediaocean). Innovid Corp. (SEC filing), 2025-02-13. https://www.sec.gov/Archives/edgar/data/1835378/000114036125004373/ef20043552_8k.htm Accessed 2026-11-10.

[743] Pixalate Announces $18.1 Million Additional Growth Capital to Accelerate Global Expansion and Media Ratings Platform. Pixalate (PR Newswire), 2021-09-02. https://www.prnewswire.com/news-releases/pixalate-announces-18-1-million-additional-growth-capital-to-accelerate-global-expansion-and-media-ratings-platform-301367707.html Accessed 2026-11-09.

[744] Star Mountain Announces Exit of Channel Factory Investment. Star Mountain Capital, 2025-06-05 (event 2025-04-24). https://starmountain.com/channel-factory/ Accessed 2026-11-09.

[745] About Fraudlogix. Fraudlogix, 2025 (event sitemap lastmod 2025-08-26). https://www.fraudlogix.com/about/ Accessed 2026-11-10.

[746] Fraud Prevention & Bot Detection (home page). Fraudlogix, 2026. https://www.fraudlogix.com/ Accessed 2026-11-10.

[747] Kubient, Inc. Form 10-K for fiscal year 2022 (Item 1, Business). Kubient, Inc. (SEC filing), 2023-03-30. https://www.sec.gov/Archives/edgar/data/1729750/000141057823000455/kbnt-20221231×10k.htm Accessed 2026-11-10. *Access: partial extract.*

[748] Kubient, Inc. Form 8-K (Items 4.01, 4.02, 5.02). Kubient, Inc. (SEC filing), 2023-09-25 (event 2023-09-19). https://www.sec.gov/Archives/edgar/data/1729750/000114036123045117/ef20011231_8k.htm Accessed 2026-11-10.

[749] EDGAR company filing index, Kubient, Inc. (CIK 0001729750). U.S. Securities and Exchange Commission (EDGAR), live index. https://www.sec.gov/cgi-bin/browse-edgar?company=kubient&type=&dateb=&owner=include&count=40&action=getcompany Accessed 2026-11-10.

[750] FAQs. Mantis (Reach plc), undated (© 2026 footer). https://www.mantissolutions.com/faqs Accessed 2026-11-10.

[751] For Advertisers & Agencies. Mantis (Reach plc; site built by Reach Solutions), undated (© 2026 footer). https://www.mantissolutions.com/for-advertisers-and-agencies Accessed 2026-11-10.

[752] Mantis Launches 'Mantis Video' to Help Publishers Compete for Video-First Briefs. Mantis (reprinting VideoWeek), 2026-06-02 (event 2026-06-02 (Mantis Video launch)). https://www.mantissolutions.com/b/mantis-launches-mantis-video-to-help-publishers-compete-for-video-first-briefs Accessed 2026-11-10.

[753] About illuma (We are illuma). Illuma Technology Limited, undated (sitemap lastmod 2026-10-15). https://weareilluma.com/about-us/ Accessed 2026-11-10.

[754] DoubleVerify Reports Third Quarter 2021 Financial Results (Exhibit 99.1). DoubleVerify Holdings, Inc. (SEC filing), 2021-11-09. https://www.sec.gov/Archives/edgar/data/1819928/000155837021015297/dv-20211109xex99d1.htm Accessed 2026-11-10.

[755] DoubleVerify Fourth Quarter and Full Year 2021 results (Exhibit 99.1). DoubleVerify Holdings, Inc. (SEC filing), 2022-03-08. https://www.sec.gov/Archives/edgar/data/1819928/000155837022003058/dv-20220308xex99d1.htm Accessed 2026-11-10.

[756] Impact Radius Acquires Forensiq, Creating the World's First Marketing System of Record with Built-In Fraud Detection. Impact (impact.com), 2016-06-14. https://impact.com/press-releases/impact-radius-acquires-forensiq-creating-the-worlds-first-marketing-system-of-record-with-built-in-fraud-detection/ Accessed 2026-11-10.

[757] Domain check: openslate.com, www.openslate.com, meetrics.com, www.meetrics.com. This stream (direct observation with curl -I), n/a. https://www.openslate.com/ Accessed 2026-11-10.

[758] Adloox Limited (08701942): company overview, officers, persons with significant control and filing history. Companies House (UK), 2026-10-01. https://find-and-update.company-information.service.gov.uk/company/08701942 Accessed 2026-11-10.

[759] Adloox Limited (08701942): Confirmation Statement made on 23 September 2025 (CS01), with shareholder details. Adloox Limited, filed at Companies House (UK), 2025-09-24. https://find-and-update.company-information.service.gov.uk/company/08701942/filing-history/MzQ4MjUxMzc0NWFkaXF6a2N4/document?format=pdf&download=0 Accessed 2026-11-10.

[760] Peer39: Our Team (company milestones). Peer39 (event 2019). https://www.peer39.com/our-team Accessed 2026-11-09.

[761] P39 TECHNOLOGIES UK LIMITED (12188698): company overview, persons with significant control and officers. Companies House (UK), 2026. https://find-and-update.company-information.service.gov.uk/company/12188698 Accessed 2026-11-10.

[762] EDGAR company record and filing index for ZEFR, INC. (CIK 0001474764; company search and data.sec.gov submissions JSON). U.S. Securities and Exchange Commission (EDGAR), 2026. https://data.sec.gov/submissi

ons/CIK0001474764.json Accessed 2026-11-10.

[763] Form D notice of exempt offering, MOVIECLIPS, INC. (accession 0001474764-12-000002). MOVIECLIPS, INC. (filed with the SEC), 2012-08-16 (event 2012-08-03). https://www.sec.gov/Archives/edgar/data/1474764/000147476412000002/xslFormDX01/primary_doc.xml Accessed 2026-11-10.

[764] Form D notice of exempt offering, ZEFR, INC. (accession 0001474764-14-000001). ZEFR, INC. (filed with the SEC), 2014-02-28 (event 2014-02-20). https://www.sec.gov/Archives/edgar/data/1474764/000147476414000001/xslFormDX01/primary_doc.xml Accessed 2026-11-10.

[765] Form D notice of exempt offering, ZEFR, INC. (accession 0001474764-14-000002). ZEFR, INC. (filed with the SEC), 2014-09-19 (event 2014-09-05). https://www.sec.gov/Archives/edgar/data/1474764/000147476414000002/xslFormDX01/primary_doc.xml Accessed 2026-11-10.

[766] Brand Safety & Suitability: Brand Safety and Ad Placement Monitoring. mFilterIt, undated. https://www.mfilterit.com/brand-safety-solution Accessed 2026-11-10.

[767] Contact Us. mFilterIt, undated. https://www.mfilterit.com/contact-us Accessed 2026-11-10.

[768] UK CMA case finder queries (gov.uk/cma-cases). Competition and Markets Authority. https://www.gov.uk/cma-cases?keywords=DoubleVerify Accessed 2026-11-10.

[769] European Commission competition case search. European Commission, DG Competition. https://competition-cases.ec.europa.eu/search?search=DoubleVerify&sortField=relevance&sortOrder=DESC Accessed 2026-11-10. *Access: partial extract.*

[770] Introducing Apostra. Apostra (formerly Scope3), 2026-09-23. https://apostra.com/blog/introducing-apostra Accessed 2026-11-09.

[771] EDGAR company filing list, DoubleVerify Holdings, Inc. (CIK 0001819928). U.S. Securities and Exchange Commission (EDGAR), 2026-11-10. https://www.sec.gov/cgi-bin/browse-edgar?action=getcompany&CIK=1819928&type=&dateb=&owner=include&count=40 Accessed 2026-11-10.

[772] DoubleVerify investor relations press-release RSS feed. DoubleVerify Holdings, Inc., 2026-11-10. https://ir.doubleverify.com/news-events/press-releases/rss Accessed 2026-11-10.

[773] 4As Cross Industry Initiatives (APB section). American Association of Advertising Agencies (4A's). https://www.aaaa.org/programs/government-relations/initiatives/4as-cross-industry-initiatives/ Accessed 2026-11-10.

[774] IAB Tech Lab Content Taxonomy. IAB Tech Lab, 2026-02-11. https://iabtechlab.com/standards/content-taxonomy/ Accessed 2026-11-10.

[775] Content Taxonomy 3.1 (TSV). IAB Tech Lab (GitHub: InteractiveAdvertisingBureau/Taxonomies), 2024-12. https://github.com/InteractiveAdvertisingBureau/Taxonomies/blob/develop/Content%20Taxonomies/Content%20Taxonomy%203.1.tsv Accessed 2026-11-10.

[776] Content Taxonomy 3.0 Descriptive Vectors (TSV). IAB Tech Lab (GitHub), 2022-06. https://github.com/InteractiveAdvertisingBureau/Taxonomies/blob/develop/Content%20Taxonomies/Content%20Taxonomy%203.0%20Descriptive%20Vectors.tsv Accessed 2026-11-10.

[777] Who We Are (JICWEBS). JICWEBS Ltd, 2019-12-06. https://web.archive.org/web/20191206054248id_/https://jicwebs.org/about-us/who-we-are/ Accessed 2026-11-10.

[778] Brand Safety (JICWEBS standards). JICWEBS Ltd, 2019-12-15. https://web.archive.org/web/20191215161359id_/https://jicwebs.org/standards/brand-safety/ Accessed 2026-11-10.

[779] How We Work (JICWEBS). JICWEBS Ltd, 2019-12-15. https://web.archive.org/web/20191215161810id_/https://jicwebs.org/about-us/how-we-work/ Accessed 2026-11-10.

[780] Ad Fraud (JICWEBS standards). JICWEBS Ltd, 2019-12-06. https://web.archive.org/web/20191206064014id_/https://jicwebs.org/standards/ad-fraud/ Accessed 2026-11-10.

[781] Viewability (JICWEBS standards). JICWEBS Ltd, 2019-12-06. https://web.archive.org/web/20191206050710id_/https://jicwebs.org/standards/viewability/ Accessed 2026-11-10.

[782] Home (JICWEBS). JICWEBS Ltd, 2019-12-15. https://web.archive.org/web/20191215161232id_/https://jicwebs.org/ Accessed 2026-11-10.

[783] TAG and JICWEBS Announce Merger of Digital Ad Industry's Leading Anti-Fraud and Brand Safety Standards Organizations. Trustworthy Accountability Group (TAG), 2020-07-29. https://www.tagtoday.net/pressreleases/tag_jicwebs_merger_announcement_7_29_2020 Accessed 2026-11-10.

[784] JICWEBS LIMITED: persons with significant control. Companies House. https://find-and-update.company-information.service.gov.uk/company/10353187/persons-with-significant-control Accessed 2026-11-10.

[785] JICWEBS LIMITED (company number 10353187): company overview. Companies House (UK registrar). https://find-and-update.company-information.service.gov.uk/company/10353187 Accessed 2026-11-10.

[786] The Content Verification Guide. AOP, IAB UK, IPA, ISBA, JICWEBS and Newsworks, 2020-01-21. https://www.isba.org.uk/system/files/media/documents/2020-12/content-verification-guide.pdf Accessed 2026-11-10.

[787] Content Verification Guide (knowledge hub page). ISBA, 2020-01-21. https://www.isba.org.uk/knowledge/content-verification-guide Accessed 2026-11-10.

[788] AAO policy registry: Common Sense Brand Standards (csbs) record and revision history. AgenticAdvertising.org, 2026-03-13 (event 2026-04-18). https://agenticadvertising.org/api/policies/resolve?policy_id=csbs Accessed 2026-11-10.

[789] Policy Registry (AdCP documentation, governance). AgenticAdvertising.org (AdCP), 2026-11-09. https://github.com/adcontextprotocol/adcp/blob/v3.2.9/docs/governance/policy-registry.mdx Accessed 2026-11-10.

[790] Brand Standards product page. Apostra, Inc., undated. https://apostra.com/solutions/brand-agents/brand-standards Accessed 2026-11-10.

[791] FTC Takes Action to Restore Competition in the Digital Advertising Ecosystem (press release). U.S. Federal Trade Commission, 2026-04-15. https://www.ftc.gov/news-events/news/press-releases/2026/04/ftc-takes-action-restore-competition-digital-advertising-ecosystem Accessed 2026-11-10.

[792] FTC et al. v. Dentsu US, Inc., GroupM Worldwide LLC (WPP Media) and Publicis, Inc., Complaint (ECF 1), 4:26-cv-00469 (N.D. Tex.). U.S. Federal Trade Commission and the States of FL, IN, IA, MT, NE, TX, UT and WV, 2026-04-15. https://storage.courtlistener.com/recap/gov.uscourts.txnd.418330/gov.uscourts.txnd.418330.1.0.pdf Accessed 2026-11-10.

[793] Responsible Media Guide (knowledge hub page). ISBA (with The Aperto Partnership), 2025-03-20 (event 2026-05 (page: "Updated May 2026")). https://www.isba.org.uk/knowledge/responsible-media-guide Accessed 2026-11-10.

[794] Martin v. Zagorski, Clerk's Judgment (ECF 62). U.S. District Court, S.D.N.Y. (Clerk of Court), 2026-09-17. https://storage.courtlistener.com/recap/gov.uscourts.nysd.639801/gov.uscourts.nysd.639801.62.0.pdf Accessed 2026-11-10.

[795] Martin v. Zagorski et al, public docket page. PacerMonitor (commercial PACER mirror), 2025-03-31. https://www.pacermonitor.com/public/case/57485217/Martin_v_Zagorski_et_al Accessed 2026-11-10.

[796] Oklahoma Firefighters v. IAS, Clerk's Judgment (ECF 63). U.S. District Court, S.D.N.Y. (Clerk of Court), 2026-03-05. https://storage.courtlistener.com/recap/gov.uscourts.nysd.635896/gov.uscourts.nysd.635896.63.0.pdf Accessed 2026-11-10.

[797] Integral Ad Science Holding Corp. (NASDAQ: IAS) Securities Fraud Class Action, case page. Kessler Topaz Meltzer & Check, LLP (plaintiffs' securities firm) (event 2026-03-26). https://www.ktmc.com/new-cases/integral-ad-science-holding-corp/ Accessed 2026-11-10.

[798] SEC EDGAR full-text search queries. U.S. Securities and Exchange Commission. https://efts.sec.gov/LATEST/search-index?q=%22Adalytics%22 Accessed 2026-11-10.

[799] EDGAR company filing index, DoubleVerify Holdings, Inc. (CIK 0001819928). U.S. Securities and Exchange Commission (EDGAR), live index. https://www.sec.gov/cgi-bin/browse-edgar?action=getcompany&CIK=0001819928&type=&dateb=&owner=include&count=20 Accessed 2026-11-10.

[800] IAS Announces Completion of Acquisition by Novacap (Exhibit 99.1). Integral Ad Science (release filed with the SEC), 2025-12-23. https://www.sec.gov/Archives/edgar/data/1842718/000114036125046475/ef20061707_ex99-1.htm Accessed 2026-11-10.

[801] Montgomery Venkersammy v. Human Security, Inc., 2:23-cv-07914 (C.D. Cal.), docket. CourtListener (Free Law Project), from PACER/RECAP, 2024-05-22. https://www.courtlistener.com/docket/67817997/montgomery-venkersammy-v-human-security-inc/ Accessed 2026-11-10. *Access: partial extract.*

[802] CourtListener RECAP search: party name "HUMAN Security". CourtListener (Free Law Project), 2026-11-10. https://www.courtlistener.com/?type=r&party_name=%22HUMAN%20Security%22&order_by=dateFiled%20desc Accessed 2026-11-10.

[803] CourtListener RECAP search: Party Name: Pixalate (5 results). CourtListener (Free Law Project), 2026-11-10. https://www.courtlistener.com/?type=r&party_name=Pixalate&order_by=dateFiled%20desc Accessed 2026-11-10.

[804] X Corp. v. World Federation of Advertisers, No. 7:24-cv-00114 (N.D. Tex.), ECF 1-2: Exhibit A to the complaint. X Corp. (plaintiff), filed in the U.S. District Court for the Northern District of Texas, 2024-08-06 (event 2022-03-16). https://storage.courtlistener.com/recap/gov.uscourts.txnd.393003/gov.uscourts.txnd.393003.1.2.pdf Accessed 2026-11-10.

[805] CourtListener RECAP search: party name 'Zefr' (all courts, all years). CourtListener (Free Law Project), 2026. https://www.courtlistener.com/?type=r&party_name=Zefr&order_by=dateFiled%20desc Accessed 2026-11-10.

[806] CourtListener RECAP full-text search: "Zefr". CourtListener (Free Law Project), 2026. https://www.courtlistener.com/?type=r&q=%22Zefr%22&order_by=dateFiled%20desc Accessed 2026-11-10. *Access: partial extract.*

[807] Accountability (AdCP documentation, media buy advanced topics). AgenticAdvertising.org (AdCP), 2026-04-08 (event 2026-10-25). https://github.com/adcontextprotocol/adcp/blob/v3.2.9/docs/media-buy/advanced-topics/accountability.mdx Accessed 2026-11-10.

[808] Billing Authority (AdCP documentation, media buy advanced topics). AgenticAdvertising.org (AdCP), 2026-05-18 (event 2026-10-29). https://github.com/adcontextprotocol/adcp/blob/v3.2.9/docs/media-buy/advanced-topics/billing-authority.mdx Accessed 2026-11-10.

[809] get_adcp_capabilities response schema, measurement block. AgenticAdvertising.org (AdCP), 2026-06-18. https://github.com/adcontextprotocol/adcp/blob/v3.2.9/static/schemas/source/protocol/get-adcp-capabilities-response.json Accessed 2026-11-10.

[810] AAO registry server source at v3.2.9: server/src/routes/registry-api.ts (listAgents OpenAPI definition) and migration 462_measurement_capabilities_snapshot.sql. AgenticAdvertising.org (AdCP repository), 2026-11-09. https://github.com/adcontextprotocol/adcp/blob/v3.2.9/server/src/routes/registry-api.ts Accessed 2026-11-10.

[811] Compliance catalog (AdCP documentation, building/verification). AgenticAdvertising.org (AdCP), 2026-11-09. https://github.com/adcontextprotocol/adcp/blob/v3.2.9/docs/building/verification/compliance-catalog.mdx Accessed 2026-11-10.

[812] Known Limitations (AdCP documentation, reference). AgenticAdvertising.org (AdCP), 2026-11-09. https://github.com/adcontextprotocol/adcp/blob/v3.2.9/docs/reference/known-limitations.mdx Accessed 2026-11-10.

[813] FAQ (AdCP documentation). AgenticAdvertising.org (AdCP), 2026-11-09. https://github.com/adcontextprotocol/adcp/blob/v3.2.9/docs/faq.mdx Accessed 2026-11-10.

[814] AAMP - Agentic Advertising Management Protocols (hub README). IAB Tech Lab (GitHub IABTechLab/AAMP), 2026-11-02. https://github.com/IABTechLab/AAMP Accessed 2026-11-10.

[815] Seller verification (AdCP documentation). AgenticAdvertising.org (AdCP), 2026-11-09. https://github.com/adcontextprotocol/adcp/blob/v3.2.9/docs/verification/overview.mdx Accessed 2026-11-10.

[816] Agentic Advertising Management Protocols from IAB Tech Lab (AAMP standards page). IAB Tech Lab, 2026-02-26 (event 2026-10-27). https://iabtechlab.com/standards/aamp-agentic-advertising-management-protocols/ Accessed 2026-11-10.

[817] DoubleVerify: ad fraud drops 41% in North America, 45% in EMEA (report on DV's 2026 Global Insights regional reports 'Media Quality in the Age of AI'). Luis Rijo, PPC Land, 2026-07-30 (event 2026-07-29 (DV's Global Insights newsletter distributing the regional reports)). https://ppc.land/doubleverify-ad-fraud-drops-41-in-north-america-45-in-emea/ Accessed 2026-11-09.

[818] Pixability home page (YouTube + CTV Data and Media Solutions). Pixability, Inc., undated (sitemap lastmod 2026-11-03). https://www.pixability.com/ Accessed 2026-11-10.

[819] Brand Suitability on YouTube (Brand Safety & Suitability, Powered by Creator Intelligence). Pixability, Inc., undated (sitemap lastmod 2026-11-03). https://www.pixability.com/brand-suitability Accessed 2026-11-10.

[820] Settlement Agreement and Mutual General Release (Exhibit 10.12 to a Phunware, Inc. filing). Phunware, Inc. (SEC EDGAR), 2020-10-09. https://www.sec.gov/Archives/edgar/data/1665300/000162828020016344/ex1012-settlementagreement.htm Accessed 2026-11-10.

[821] Publishers are underwhelmed by the payoff from hitting viewability standards. Ross Benes, Digiday, 2017-12-20. https://digiday.com/media/hitting-viewability-thresholds-pubs-underwhelmed-payoff/ Accessed 2026-11-10.

[822] EDGAR company filing list, DoubleVerify Holdings, Inc. (CIK 0001819928). U.S. Securities and Exchange Commission (EDGAR), 2026-11-10. https://www.sec.gov/cgi-bin/browse-edgar?action=getcompany&CIK=1819928&type=&dateb=&owner=include&count=10 Accessed 2026-11-10.

[823] Sightly home page. Sightly Enterprises, Inc., undated (page footer: (c) 2026). https://www.sightly.com/ Accessed 2026-11-10.

[824] Accreditation letter to Flashtalking (now New Innovid): device reports incl. Protected Media-measured traffic (GI030525). Media Rating Council (MRC), 2025-03-05. https://mediaratingcouncil.org/sites/default/files/Accreditation%20Letters/Flashtalking/GI030525_Flashtalking%20Accreditation%20Letter.pdf Accessed 2026-11-10.

[825] HUMAN Documentation home and llms.txt documentation index. HUMAN Security (documentation site), undated. https://docs.humansecurity.com/llms.txt Accessed 2026-11-10.

[826] FouAnalytics (home page: Self Serve request-invite form, Sign In, Humanness/hCPM, PageXray, Anonymously Report Fraud, links). FouAnalytics (Marketing Science Consulting Group, Inc.), undated. https://www.fouanalytics.com/ Accessed 2026-11-10.

[827] Media Quality (solution page). Zefr, undated. https://zefr.com/solution/media-quality Accessed 2026-11-10.

[828] Apostra home page. Apostra, Inc., undated (Copyright © 2026 Apostra, Inc.). https://apostra.com/ Accessed 2026-11-10.

[829] NewsGuard for Advertising. NewsGuard, undated. https://www.newsguardtech.com/solutions/newsguard-for-advertising/ Accessed 2026-11-10.

[830] IRIS.TV home page (Video Data Connectivity). IRIS.TV, undated. https://www.iris.tv/ Accessed 2026-11-10.

[831] Vobile (us.vobile.com). Vobile Group Limited, undated. https://us.vobile.com/ Accessed 2026-11-10.

[832] Click Fraud Protection Pricing for Every Business Size. TrafficGuard (Adveritas), undated (sitemap lastmod 2026-10-29). https://www.trafficguard.ai/pricing Accessed 2026-11-10.

[833] Ad Tag (help-centre article). TrafficGuard Support, 2023-01-03 (date as rendered via curl; 2 January in the browser per the 9 Nov check). https://help.trafficguard.ai/en/articles/6751345-ad-tag Accessed 2026-11-10.

[834] Measurement API (help-centre article). TrafficGuard Support, n/a. https://help.trafficguard.ai/en/articles/6751354-measurement-api Accessed 2026-11-10. *Access: partial extract.*

[835] TrafficGuard sitemap.xml (668 URLs). TrafficGuard (Adveritas), lastmod values to 2026-11-03. https://www.trafficguard.ai/sitemap.xml Accessed 2026-11-10.

[836] Click Fraud Protection & Prevention for Advertisers (home page). TrafficGuard, undated (©2026 TrafficGuard). https://www.trafficguard.ai/ Accessed 2026-11-10.

[837] TrafficGuard expands New York base amid US ad fraud surge. TrafficGuard, 2026-02-18. https://www.trafficguard.ai/news/trafficguard-expands-new-york-base-amid-us-ad-fraud-surge Accessed 2026-11-10.

[838] CHEQ home page (Identity Security for the Agentic Web). CHEQ AI Technologies Ltd., undated (© 2026). https://cheq.ai/ Accessed 2026-11-10.

[839] Verify and protect your display ad traffic. Lunio, undated (Lunio 2026). https://www.lunio.ai/solutions/display-ads Accessed 2026-11-10.

[840] Spider AF home page (Marketing Fraud Prevention Platform). Spider AF, undated. https://spideraf.com/ Accessed 2026-11-10.

[841] fraud0 home page (Recover wasted ad spend). fraud0, undated (© 2026 fraud0). https://www.fraud0.com/ Accessed 2026-11-10.

[842] Ad Fraud Protection Solution. AppsFlyer, undated. https://www.appsflyer.com/products/measurement/fraud-protection/ Accessed 2026-11-10.

[843] GeoEdge home page. GeoEdge, undated. https://www.geoedge.com/ Accessed 2026-11-10.

[844] Confiant home page. Confiant, undated. https://www.confiant.com/ Accessed 2026-11-10.

[845] The Media Trust home page. The Media Trust, undated. https://mediatrust.com/ Accessed 2026-11-10.

[846] xpln.ai home page. xpln.ai, undated. https://xpln.ai/ Accessed 2026-11-10.

[847] Case Study: Numberly boosts conversion rate by 10% using pre-bid attention segments. Adloox, 2024 (event 2024-01-08 to 2024-01-25). https://www.adloox.com/customers/case-studies/numberly Accessed 2026-11-10.

[848] OpenAlex works search 'adloox' (11 results) and arXiv API search 'adloox' (0 results). OpenAlex (OurResearch); arXiv (Cornell University), 2026. https://api.openalex.org/works?search=adloox Accessed 2026-11-10. *Access: abstract only.*

[849] DV Uncovers New Large-Scale Audio Ad Fraud Scheme that Spoofed Over 500,000 Devices in March 2024 (release). DoubleVerify, 2024-05-07. https://doubleverify.com/company/newsroom/dv-uncovers-new-large-scale-audio-ad-fraud-scheme-that-spoofed-over-500000-devices-in-march-2024 Accessed 2026-11-10.

[850] DoubleVerify Expands Measurement Coverage on TikTok Pangle In the United States and International Markets (release, 13 August 2026). DoubleVerify Holdings, Inc. (GlobeNewswire release on DV investor relations), 2026-08-13. https://ir.doubleverify.com/news-events/press-releases/press-releases-details/2026/DoubleVerify-Expands-Measurement-Coverage-on-TikTok-Pangle-In-the-United-States-and-International-Markets/default.aspx Accessed 2026-11-10.

[851] IAS Help Center: Brand Safety - Risk Bucket Thresholds (login redirect). Integral Ad Science. https://helpcenter.integralplatform.com/article/brand-safety-risk-buckets Accessed 2026-11-10. *Access: partial extract.*

[852] Social Platforms solutions page. Integral Ad Science, 2025-04-24. https://integralads.com/solutions/social-platforms/ Accessed 2026-11-10.

[853] Full-Funnel Ad Traffic Validation. mFilterIt, undated. https://www.mfilterit.com/full-funnel-ad-traffic-validation Accessed 2026-11-10.

[854] Is Your Ad Fraud Verification Partner Using the Latest Technology?. mFilterIt (mFilterIt Experts), 2024-12-03. https://www.mfilterit.com/blog/ad-fraud-verification-partner-using-latest-technology Accessed 2026-11-10.

[855] One Tag for All: The New Tech in Bot Detection to Protect Your Ad Campaigns in 2026. mFilterIt, 2026-08-13. https://www.mfilterit.com/blog/bot-detection-one-tag-ad-fraud Accessed 2026-11-10.

[856] Contact mFilterIt USA | Digital Trust Experts. mFilterIt (MFILTERIT TECHNOLOGIES INC), undated. https://www.mfilterit.com/contact-us-usa Accessed 2026-11-10.

[857] Ad Fraud Report: Beyond the Linear Lens (report landing page). mFilterIt, 2025-12-09. https://www.mfilterit.com/reports/ad-fraud-report-2025 Accessed 2026-11-10.

[858] Beyond Brand Safety: Why YouTube Relevancy Analysis Matters for Advertisers. mFilterIt (mFilterIt Experts), 2026-11-03. https://www.mfilterit.com/blog/youtube-brand-safety-relevancy-analysis Accessed 2026-11-10.

[859] CTV Scale integrates mFilterIt to enhance inventory quality and trust across Connected TV ecosystem. MediaNews4U (MN4U Bureau), 2026-01-07 (event 2026-01-06 (earliest copy found, marketingmind.in)). https://www.medianews4u.com/ctv-scale-integrates-mfilterit-to-enhance-inventory-quality-and-trust-across-connected-tv-ecosystem/ Accessed 2026-11-10.

[860] Meta Ad Fraud Detection. mFilterIt, undated. https://www.mfilterit.com/meta-ads-fraud-detection Accessed 2026-11-10.

[861] About mFilterIt. mFilterIt, undated. https://www.mfilterit.com/about-us Accessed 2026-11-10.

[862] Neutrality is Key to Third-Party Audits and Validation. mFilterIt (Jyoti Kalra), 2021-02-25. https://www.mfilterit.com/blog/neutrality-is-key-to-third-party-audits-and-validation Accessed 2026-11-10.

[863] mPulse by mFilterIt: One Agentic Operating Layer for Every Media Campaign. mFilterIt, undated. https://www.mfilterit.com/agentic-ai-marketing-platform-branding-campaigns Accessed 2026-11-10.

[864] Brand Suitability | Protected by Mediaocean. Protected by Mediaocean, undated (footer 2026). https://www.protected.media/brand-suitability Accessed 2026-11-10.

[865] Viewability & Attention | Protected by Mediaocean. Protected by Mediaocean, undated (footer 2026). https://www.protected.media/viewability-attention Accessed 2026-11-10.

[866] Privacy Policy (Protected Media Ltd.). Protected Media Ltd., 2026-09-11. https://www.protected.media/privacy-policy Accessed 2026-11-10.

[867] The Infrastructure of Understanding: Inclusion Over Exclusion (blog, Denys Dantsev). Protected by Mediaocean, 2026-01-21. https://www.protected.media/infrastructure-understanding-inclusion-over-exclusion Accessed 2026-11-10.

[868] Identifying wasteful inventory & trends (Princess Cruises case study). Protected by Mediaocean, 2025-10-20. https://www.protected.media/case-study-princess-wasteful-inventory Accessed 2026-11-10.

[869] Path to premium ad quality (UScellular case study). Protected by Mediaocean, 2025-10-20. https://www.protected.media/case-study-uscellular-ad-quality Accessed 2026-11-10.

[870] Operation Slay Hydra Advisory (landing page). Protected by Mediaocean, 2025-01-15. https://www.protected.media/operation-slay-hydra-advisory Accessed 2026-11-10. *Access: partial extract.*

[871] Basis Integrates Protected by Mediaocean to Deliver AI-Driven, Real-Time Media Quality & Performance Signals into Campaign Activation. Basis and Mediaocean (joint release on protected.media), 2026-04-14. https://www.protected.media/protected-by-mediaocean-basis-integration Accessed 2026-11-10.

[872] Why Pre-Bid Ad Verification Needs a Feedback Loop (blog, Sara Maskivish). Protected by Mediaocean, 2026-08-20. https://www.protected.media/why-pre-bid-ad-verification-needs-feedback-loop Accessed 2026-11-10.

[873] Solutions | Protected by Mediaocean. Protected by Mediaocean, undated (footer 2026). https://www.protected.media/solutions Accessed 2026-11-10.

[874] Protected Media Launches ProtectedTV: A Technology Suite for End-to-end OTT Ad Fraud Protection. Protected Media (release), 2020-08-20. https://www.protected.media/protected-media-launches-protectedtv-a-technology-suite-for-end-to-end-ott-ad-fraud-protection Accessed 2026-11-10.

[875] Protected Media Launches OTT Supply Side Ad Verification. Protected Media (release), 2019-09-24. https://www.protected.media/ott-supply-side-verification Accessed 2026-11-10.

[876] New Brand Safety and Suitability Control for Ads on Threads (Meta for Business announcement). Meta, 2026-05-18. https://www.facebook.com/business/news/brandsafety Accessed 2026-11-10.

[877] The Era of Understanding: A Smarter Approach to Brand Safety and Suitability on Meta (blog). Protected by Mediaocean, 2026-05-28. https://www.protected.media/the-era-of-understanding-a-smarter-approach-to-brand-safety-and-suitability-on-meta Accessed 2026-11-10.

[878] Protected partners with Internet Watch Foundation to strengthen digital safety. Protected by Mediaocean (release), 2025-08-14. https://www.protected.media/protected-partners-with-internet-watch-foundation-to-strengthen-digital-safety Accessed 2026-11-10.

[879] Prisma Integration | Protected by Mediaocean. Protected by Mediaocean, undated (footer 2026). https://www.protected.media/solutions/prisma-integration Accessed 2026-11-10.

[880] Terms of Use (Protected Media Ltd.). Protected Media Ltd., 2026-09-15. https://www.protected.media/terms-of-use Accessed 2026-11-10.

[881] July Release Notes (blog). Protected by Mediaocean, 2025-07-21. https://www.protected.media/protected-by-mediaocean-july-release-notes Accessed 2026-11-10.

[882] Ad Viewability. Escalated Solutions, LLC, undated. https://escalated.io/ad-viewability.html Accessed 2026-11-10.

[883] Escalated.io Releases Next Generation Software Update For Its Invalid Traffic Filtration Service. Escalated Solutions, LLC, 2026-03-16. https://escalated.io/n/Escalated.io-Releases-Next-Generation-Software-Update-For-Its-Invalid-Traffic-Filtration-Service.html Accessed 2026-11-10.

[884] Case Study: Fixing a Billion-Dollar Ad Fraud Problem: How Escalated Empowered PPCmate with Real-Time Traffic Protection. Escalated Solutions, LLC, undated. https://escalated.io/case-studies/case-study-1.html Accessed 2026-11-10.

[885] Case Study: FirstImpression.io Enhances News Publisher Success with Escalated.io's IVT Detection. Escalated Solutions, LLC, undated. https://escalated.io/case-studies/case-study-2.html Accessed 2026-11-10.

[886] Real-Time Ad Fraud Protection (partners and FAQ). PPCmate (3DMA Ltd), 2024-11-15 (datePublished; dateModified 2026-11-10). https://ppcmate.com/real-time-ad-protection/ Accessed 2026-11-10.

[887] IVT Mitigation: Ad Fraud Prevention & Monitoring. Escalated Solutions, LLC, undated. https://escalated.io/ad-fraud.html Accessed 2026-11-10.

[888] IVT Integration Methods. Escalated Solutions, LLC, undated. https://escalated.io/integrations.html Accessed 2026-11-10.

[889] Contact Us. Escalated Solutions, LLC, undated. https://escalated.io/contact-us.html Accessed 2026-11-10.

[890] Privacy (served at escalated.io/privacy.html from the fragment privacy.fragment). Escalated Solutions, LLC, undated. https://escalated.io/privacy.fragment Accessed 2026-11-10.

[891] Terms of Service (served at escalated.io/terms.html from the fragment terms.fragment). Escalated Solutions, LLC, undated. https://escalated.io/terms.fragment Accessed 2026-11-10.

[892] Interviewed With Kyle Smith - Founder and CEO of Escalated. SafetyDetectives (Shauli Zacks, content editor), 2025-10-15. https://www.safetydetectives.com/blog/kyle-smith-escalated/ Accessed 2026-11-10.

[893] Cybersecurity Measures Built Into FouAnalytics. Dr. Augustine Fou (founder), 2020-01-26. https://www.linkedin.com/pulse/does-your-analytics-platform-have-cybersecurity-dr-augustine-fou- Accessed 2026-11-10.

[894] Why It's Hard for Bots To Avoid FouAnalytics Detection. Dr. Augustine Fou (founder), 2022-01-02. https://www.linkedin.com/pulse/why-its-hard-bots-avoid-fouanalytics-detection-the-independent/ Accessed 2026-11-10.

[895] How FouAnalytics is different from DoubleVerify and IAS. Dr. Augustine Fou (founder), 2025-03-27. https://www.linkedin.com/pulse/how-fouanalytics-compares-doubleverify-ias-dr-augustine-fou-igque/ Accessed 2026-11-10.

[896] Tour of FouAnalytics Dashboard and Campaign Checklist. Dr. Augustine Fou (founder), 2022-11-19. https://www.linkedin.com/pulse/tour-fouanalytics-dashboard-campaign-checklist-dr-augustine Accessed 2026-11-10.

[897] FouAnalytics Announces Global Availability of Premium Feature -- Content-Level Brand Safety Analytics With Full Page-URL Reporting. FouAnalytics (PR Newswire), 2026-08-19. https://www.prnewswire.com/news-releases/fouanalytics-announces-global-availability-of-premium-feature----content-level-brand-safety-analytics-with-full-page-url-reporting-302855543.html Accessed 2026-11-10.

[898] FouAnalytics announces global availability of FouAnalytics Unlimited (corrected release). FouAnalytics (PR Newswire), 2026-08-06 (event 2026-08-07 (correction: an image added)). https://www.prnewswire.com/news-releases/fouanalytics-announces-global-availability-of-fouanalytics-unlimited-302845458.html Accessed 2026-11-10.

[899] How to Deploy FouAnalytics Tags on Ad Platforms and Sites. Dr. Augustine Fou (founder), 2022-11-07. https://www.linkedin.com/pulse/how-deploy-fouanalytics-tags-various-ad-platforms-dr-augustine Accessed 2026-11-10.

[900] How to Use the Domain App Report in FouAnalytics to Review Sites and Apps. Dr. Augustine Fou (founder), 2022-10-29. https://www.linkedin.com/pulse/how-use-domain-app-report-fouanalytics-review-sites-dr-augustine Accessed 2026-11-10.

[901] FouAnalytics Announces Global Availability of FouAnalytics for Publishers, Decisioning Container. FouAnalytics (PR Newswire), 2026-08-17. https://www.prnewswire.com/news-releases/fouanalytics-announces-global-availability-of-fouanalytics-for-publishers-decisioning-container-302853144.html Accessed 2026-11-10.

[902] FouAnalytics Releases Public MSA Boilerplate for Measurement-First Digital Media Governance. FouAnalytics (PR Newswire), 2026-08-18. https://www.prnewswire.com/news-releases/fouanalytics-releases-public-msa-boilerplate-for-measurement-first-digital-media-governance-302854430.html Accessed 2026-11-10.

[903] Digital Media Governance, Verification, Transparency, and Brand Safety Master Services Agreement BOILERPLATE (Digital_Media_Governance_MSA_WORD.docx). FouAnalytics (Google Docs file linked from the release), 2026-08 (document undated; released 2026-08-18). https://docs.google.com/document/d/1O9bdRbJEYIfm_u4NZsOx2wz13-5ppW-g/ Accessed 2026-11-10.

[904] IntentKey Integrates FouAnalytics Verification to Help Advertisers Improve Programmatic Efficiency. Inuvo, Inc. (GlobeNewswire, posted on inuvo.com), 2026-08-26. https://inuvo.com/intentkey-integrates-fouanalytics-verification-to-help-advertisers-improve-programmatic-efficiency/ Accessed 2026-11-10.

[905] FouAnalytics and fraud0 Partner to Bring Automated Google Ads Refunds to the World's Largest Advertisers. FouAnalytics (PR Newswire), 2026-10-07. https://www.prnewswire.com/news-releases/fouanalytics-and-fraud0-partner-to-bring-automated-google-ads-refunds-to-the-worlds-largest-advertisers-302901699.html Accessed 2026-11-10.

[906] FouAnalytics announces global availability of in-situ screenshots for proper governance of digital campaigns. FouAnalytics (PR Newswire), 2026-08-14. https://www.prnewswire.com/news-releases/fouanalytics-announces-global-availability-of-in-situ-screenshots-for-proper-governance-of-digital-campaigns-302852060.html Accessed 2026-11-10.

[907] SEC EDGAR full-text search: "FouAnalytics" (2020-01-01 to 2026-11-10) and "Fou Analytics" (2015-01-01 to 2026-11-10). US Securities and Exchange Commission (EDGAR FTS), live search. https://efts.sec.gov/LATEST/search-index?q=%22FouAnalytics%22&dateRange=custom&startdt=2020-01-01&enddt=2026-11-10 Accessed 2026-11-10.

[908] Privacy Policy v2026. FouAnalytics / Marketing Science Consulting Group, Inc., 2026. https://fouanalytics.com/privacy Accessed 2026-11-10.

[909] SimpleFeed and FouAnalytics partner to bring independent verification to Apple News campaigns. FouAnalytics and SimpleFeed, Inc. (PR Newswire; source FouAnalytics), 2026-08-10. https://www.prnewswire.com/news-releases/simplefeed-and-fouanalytics-partner-to-bring-independent-verification-to-apple-news-campaigns-302847324.html Accessed 2026-11-10.

[910] Getting Started With FouAnalytics. Dr. Augustine Fou (founder), 2023-09-09. https://www.linkedin.com/pulse/how-start-using-fouanalytics-dr-augustine-fou/ Accessed 2026-11-10.

[911] Methodology. Fraudlogix, 2021-04-15 (event modified 2022-09-08). https://www.fraudlogix.com/how/ Accessed 2026-11-10.

[912] What is Viewability? (glossary). Fraudlogix, 2025-12-23 (event 'Last updated: Dec 2025'). https://www.fraudlogix.com/glossary/what-is-viewability/ Accessed 2026-11-10.

[913] The State of Ad Fraud 2026 (Ad Fraud Statistics 2026: 20.64% IVT Rate). Fraudlogix, 2026-01-14 (event modified 2026-11-05). https://www.fraudlogix.com/stats/ad-fraud-statistics-2026 Accessed 2026-11-10.

[914] United States v. Zhukov, No. 1:18-cr-00633-EK (E.D.N.Y.), Document 405: government sentencing letter. U.S. Attorney, Eastern District of New York, 2021-10-05 (event letter dated 2021-10-04; conduct September 2014 to December 2016). https://storage.courtlistener.com/recap/gov.uscourts.nyed.425809/gov.uscourts.nyed.425809.405.0_2.pdf Accessed 2026-11-10.

[915] IP Blocklist API Documentation. Fraudlogix, 2025-12-08 (event modified 2026-11-05). https://www.fraudlogix.com/ip-blocklist-api-documentation/ Accessed 2026-11-10.

[916] Ad Fraud Statistics & Research (stats hub). Fraudlogix, 2026 (event sitemap lastmod 2026-07-28). https://www.fraudlogix.com/stats Accessed 2026-11-10.

[917] Best Ad Fraud Solution 2026 | Detect & Block Invalid Traffic (blog). Fraudlogix (Jessica Lynn Lopez, co-founder and COO), 2026-01-08 (event 'Updated January 8, 2026'; modified 2026-11-06). https://www.fraudlogix.com/blog/best-ad-fraud-solution-2026 Accessed 2026-11-10.

[918] Advertisers are unknowingly wasting 11.5% of their budgets (ad blocker report). Fraudlogix, 2020-06-30 (event modified 2025-08-13). https://www.fraudlogix.com/adblocker-report/ Accessed 2026-11-10.

[919] AdKernel Integration. Fraudlogix, 2019-04-24 (event modified 2022-03-07). https://www.fraudlogix.com/adkernel-integration/ Accessed 2026-11-10.

[920] AdKernel Integrates Fraudlogix's Ad Verification Technology Directly into RTB Platform To Improve Advertisers' Fraud Management Capabilities. AdKernel (via ACCESSWIRE), 2019-05-29. https://www.accesswire.com/newsroom/en/business-and-professional-services/adkernel-integrates-fraudlogixs-ad-verification-technology-direct-547095 Accessed 2026-11-10.

[921] Fraudlogix Terms of Service. Fraudlogix, undated. https://ws.fraudlogix.com/terms.php Accessed 2026-11-10.

[922] Start Free Post-Bid Analytics (sign-up page). Fraudlogix, Inc., 2026. https://ws.fraudlogix.com/lite_coinfo.php Accessed 2026-11-10.

[923] Privacy Policy. Fraudlogix, 2026-10-01 (event 'Last updated and effective as of October 1, 2026'; page modified 2026-10-23). https://www.fraudlogix.com/privacy-policy/ Accessed 2026-11-10.

[924] Fraudlogix vs. DoubleVerify: Ad Fraud Detection Compared. Fraudlogix, 2026-04-24 (event modified 2026-11-05). https://www.fraudlogix.com/compare/fraudlogix-vs-doubleverify/ Accessed 2026-11-10.

[925] Fraudlogix adds reason codes to IP blocklist solution. Fraudlogix (blog.fraudlogix.com), 2017-03-09. https://blog.fraudlogix.com/fraudlogix-adds-reason-codes-to-ip-blacklist-solution Accessed 2026-11-10.

[926] Snowflake Data Fields (Advertising Protection documentation). HUMAN Security, 2026-10-27. https://docs.humansecurity.com/advertising/snowflake-data-fields Accessed 2026-11-10.

[927] Satori Threat Intelligence Alert: Apollo Spoofs Audio to Maximize Profits (HUMAN blog, 17 April 2025), Internet Archive capture of 13 May 2025. HUMAN Security (Satori Threat Intelligence and Research Team; one researcher credited from The Trade Desk), 2025-04-17. https://web.archive.org/web/20250513122253/https://www.humansecurity.com/learn/blog/satori-threat-intelligence-alert-apollo-spoofs-audio-to-maximize-profits/ Accessed 2026-11-10.

[928] Ad fraud in audio: The promise and challenges of audio advertising (HUMAN blog, 7 October 2024), Internet Archive capture of 11 July 2025. HUMAN Security, 2024-10-07. https://web.archive.org/web/20250711213731/https://www.humansecurity.com/learn/blog/ad-fraud-in-audio-thepromise-andchallenge-of-audioadvertising/ Accessed 2026-11-10.

[929] Audio Advertising Fraud: What You Need to Know (HUMAN blog, 15 November 2023), Internet Archive capture of 29 April 2025. HUMAN Security, 2023-11-15. https://web.archive.org/web/20250429111147/https://www.humansecurity.com/learn/blog/audio-advertising-fraud-what-you-need-to-know/ Accessed 2026-11-10.

[930] Nuanced Intelligence: Addressing Made for Advertising Sites (HUMAN blog, 13 March 2025), Internet Archive capture of 7 October 2025. HUMAN Security, 2025-03-13. https://web.archive.org/web/20251007052735/https://www.humansecurity.com/learn/blog/addressing-made-for-advertising-sites/ Accessed 2026-11-10.

[931] HUMAN Security: How HUMAN Security and LinkedIn work together (LinkedIn Marketing Solutions partner directory). LinkedIn Corporation, 2025. https://business.linkedin.com/marketing-solutions/marketing-partners/find-a-partner/human-security Accessed 2026-11-10.

[932] About the MediaGuard API (Advertising Protection documentation). HUMAN Security, 2026-08-20. https://docs.humansecurity.com/advertising/reference/about-the-mediaguard-api Accessed 2026-11-10.

[933] Pixalate Releases Q2 2026 North America Invalid Traffic (IVT) & Ad Fraud Benchmarks: United States Mobile Apps See 39% IVT, CTV at 25%. Pixalate, 2026-08-12. https://www.pixalate.com/blog/q2-2026-ad-fraud-be

nchmarks-report-for-north-america Accessed 2026-11-10.

[934] Pixalate and Pangle, Ad Network of TikTok for Business, Announce Global Partnership (republication of a GlobeNewswire release). Pixalate, Inc. (republication), 2025-05-19. https://www.pixalate.com/blog/pixalate-and-pangle-announce-global-partnership Accessed 2026-11-10.

[935] Gravite partners with Pixalate to raise the bar on mobile ad supply quality (republication). Pixalate, Inc. (republication of Gravite's blog), 2026-11-05. https://www.pixalate.com/blog/gravite-partners-with-pixalate-to-raise-the-bar-on-mobile-ad-supply-quality Accessed 2026-11-10.

[936] Pixalate documentation: Traffic tags in your ad server (Google Ad Manager, The Trade Desk, Beeswax, Campaign Manager 360). Pixalate, Inc.. https://docs.pixalate.com/analytics/ad-server-trafficking Accessed 2026-11-10.

[937] Screencore and Pixalate Build on Long-Standing Partnership for Ad Fraud Analytics and Traffic Quality Measurement. Pixalate, Inc. (partner release), 2026-10-22. https://www.pixalate.com/blog/screencore-and-pixalate-build-on-long-standing-partnership-for-ad-fraud-analytics-and-traffic-quality-measurement Accessed 2026-11-10.

[938] Pixalate documentation changelog entry of February 5, 2026. Pixalate, Inc., 2026-02-05. https://docs.pixalate.com/changelog/2026/2/5.md Accessed 2026-11-10.

[939] Channel Factory launches AI Slop Detection. Advanced Television, 2026-07-16. https://www.advanced-television.com/2026/07/16/channel-factory-launches-ai-slop-detection/ Accessed 2026-11-10.

[940] Channel Factory is THE Media Wastage Solution. ExchangeWire PressBox (Channel Factory announcement), 2023-05-03. https://www.exchangewire.com/blog/2023/05/03/channel-factory-is-the-media-wastage-solution/ Accessed 2026-11-10.

[941] Channel Factory's ViewIQ reduces media wastage by 23pc leading to a 12pc cost savings: PwC. MediaNews4U (MN4U Bureau), 2023-05-08. https://www.medianews4u.com/channel-factorys-viewiq-reduces-media-wastage-by-23pc-leading-to-a-12pc-cost-savings-pwc/ Accessed 2026-11-10.

[942] AIMM Growth Fronts member submission: Channel Factory. ANA AIMM Growth Fronts (profile submitted by Channel Factory). https://aimmgrowthfronts.com/dotm-submissions/channel-factory/ Accessed 2026-11-10.

[943] ViewIQ 'Inspect' screen mock-up (product page image asset 'Report Details 1'). Channel Factory. https://cdn.prod.website-files.com/67c1942c4bbd900545de0c79/67cae57acc207f10be080c94_Report%20Details%201.webp Accessed 2026-11-10.

[944] AERO landing page FAQ answers (Framer page module). Channel Factory. https://framerusercontent.com/sites/DoXePVTIVWbvIwdxw1fKG/GADz-wwBv86EuNh_kdmEdpWs91ySLfz2IO9qIzpp6IM.SoVDrr3Z.mjs Accessed 2026-11-10.

[945] Brand safety and suitability (hub). Meta for Business. https://www.facebook.com/business/brand-safety-and-suitability Accessed 2026-11-10.

[946] Google Enterprise Marketing Portal script bundle (Advertising Technology Partners directory). Google. https://www.gstatic.com/alkali/04204e8602008a7314a5ed86b069cc5300a44f86.js Accessed 2026-11-10.

[947] Success Stories (case-study listing). Channel Factory, 2026. https://channelfactory.com/insights/success-stories Accessed 2026-11-10.

[948] News UK partners with illuma to boost ad accuracy. Illuma Technology Limited, 2024-08-13. https://weareilluma.com/news-uk-and-illuma-partner-to-boost-contextual-accuracy-for-targeting-and-brand-safety/ Accessed 2026-11-10.

[949] illuma Precision™ | Human-level Contextual Intelligence. Illuma Technology Limited, undated (sitemap lastmod 2026-10-30). https://weareilluma.com/illuma-precision/ Accessed 2026-11-10.

[950] TAG Registry entry: Illuma (registry widget at tagtoday.net/registry, served from hubgloo.opxar.com). Trustworthy Accountability Group (TAG), live registry. https://hubgloo.opxar.com/company/17008483195 Accessed 2026-11-10.

[951] Drive more valuable reach across the open internet with Kokai. The Trade Desk, 2024-05-28. https://www.thetradedesk.com/resources/drive-more-valuable-reach-across-the-open-internet-kokai Accessed 2026-11-

10.

[952] illuma appoints Michel Guillon to lead solutions architecture. Illuma Technology Limited, 2026-10-06. https://weareilluma.com/illuma-appoints-michel-guillon-solutions-architecture/ Accessed 2026-11-10.

[953] illuma brings contextual-AI to Amazon's APS marketplace. Illuma Technology Limited, 2024-08-21. https://weareilluma.com/illumas-ai-powered-contextual-tools-are-now-available-through-amazon-publisher-services-connections-marketplace/ Accessed 2026-11-10.

[954] Illuma Technology Ltd: accounts for a small company made up to 31 December 2024 (iXBRL). Illuma Technology Ltd (audited by Gerald Edelman LLP); filed at Companies House, 2025-09-19 (filed; audit report dated 2025-09-18). https://find-and-update.company-information.service.gov.uk/company/06727425/filing-history/MzQ4MTk1NTE5MGFkaXF6a2N4/document?format=xhtml&download=0 Accessed 2026-11-10.

[955] Enhancing brand-audience experiences through contextual advertising (report). Illuma Technology Limited, 2023-06-15 (PDF creation date). https://weareilluma.com/wp-content/uploads/2026/06/illuma-brand-audience.2023.pdf Accessed 2026-11-10.

[956] Get in touch (contact page). Illuma Technology Limited, undated (sitemap lastmod 2026-11-02). https://weareilluma.com/contact-us/ Accessed 2026-11-10.

[957] End User Agreement and Privacy Policy. Illuma Technology Ltd, 2025-08. https://weareilluma.com/end-user-policy/ Accessed 2026-11-10.

[958] Pre-bid programmatic advertising resource guide (blog). Peer39, 2021-10-08. https://www.peer39.com/blog/pre-bid-programmatic-resource-guide Accessed 2026-11-10.

[959] Announcing New CTV Suitability Categories for Pre-bid Contextual Targeting (blog). David Simutis, Peer39, 2021-11-15. https://www.peer39.com/blog/ctv-brand-safety-suitability Accessed 2026-11-10.

[960] Made-for-Advertising (MFA) Sites: How to Avoid MFA. Peer39, undated (footer © 2026). https://www.peer39.com/made-for-advertising-sites Accessed 2026-11-10.

[961] Crawler Notice (Last updated: June 1st 2023). Peer39, 2023-06-01. https://www.peer39.com/crawler-notice Accessed 2026-11-10.

[962] Advertisers and Agencies (with FAQ). Peer39, undated (footer © 2026). https://www.peer39.com/advertisers-and-agencies Accessed 2026-11-10.

[963] Peer39's Brand Suitability and Safety Controls: How To Think About Brand Protection (blog, tagged Case Study). Dannalyn Prado, Peer39, 2025-10-02. https://www.peer39.com/blog/peer39s-brand-suitability-and-safety-controls-how-to-think-about-brand-protection Accessed 2026-11-10.

[964] Quarterly CTV benchmarks (CTV Advertising Benchmarks Report, Q3 2026). Peer39, 2026-10. https://www.peer39.com/quarterly-ctv-benchmarks Accessed 2026-11-10.

[965] Brand Suitability & Safety Controls ("Categories for CTV, Mobile, Display, OLV, and Audio"). Peer39, undated (footer © 2026). https://www.peer39.com/suitability/ Accessed 2026-11-10.

[966] Amazon DSP Advertising. Peer39, undated (footer © 2020-2026). https://www.peer39.com/amazon-dsp-solutions Accessed 2026-11-10.

[967] Traffic Quality Targeting (Zeta DSP Knowledge Base). Zeta Global, 2025-12-16 (page last-modified metadata). https://knowledgebase.zetaglobal.com/pug/traffic-quality-targeting Accessed 2026-11-10.

[968] Brand safety and media quality (Adobe Advertising DSP Guide). Adobe, 2026-10-02. https://experienceleague.adobe.com/en/docs/advertising/dsp/introduction/features/brand-safety-media-quality Accessed 2026-11-10.

[969] Content Categories Targeting (Zeta DSP Knowledge Base). Zeta Global, 2023-02-10 (page last-modified metadata). https://knowledgebase.zetaglobal.com/pug/content-categories Accessed 2026-11-10.

[970] Peer39 Primes You for Amazon DSP Success (blog). David Simutis, Peer39, 2025-06-11. https://www.peer39.com/blog/peer39-on-amazon-dsp Accessed 2026-11-10.

[971] CTV Advertising & Contextual Targeting Platform. Peer39, undated (footer © 2026). https://www.peer39.com/platform Accessed 2026-11-10.

[972] Contextual Advertising | Contextual Intelligence (home page). Peer39, undated (footer © 2026). https://www.peer39.com/ Accessed 2026-11-10.

[973] Turning CTV constraints into a 25% efficiency gain (John Moore Services case study). Peer39, undated. https://www.peer39.com/john-moore-service-home-services Accessed 2026-11-10.

[974] CTV Publisher Solutions. Peer39, undated (footer © 2020-2026). https://www.peer39.com/ctv-publisher-solutions Accessed 2026-11-10.

[975] Terms and Conditions - Amazon Connections Marketplace. Peer39 (P39 Tech, LLC), undated. https://www.peer39.com/terms-and-conditions-amazon-marketplace Accessed 2026-11-10.

[976] Pixability Announces the Launch of GenAI Contextual Segments (GCS) for YouTube. Pixability, Inc., 2024-10-07. https://www.pixability.com/about/news/launch-of-genai-contextual-segments-for-youtube-to-drive-better-targeting-and-outcomes-for-advertisers Accessed 2026-11-10.

[977] Pixability Launches Industry's First MCP-Enabled AI Agent Layer for YouTube Advertising and Creator Intelligence. Pixability, Inc., 2026-06-17. https://www.pixability.com/about/news/pixability-launches-industrys-first-mcp-enabled-ai-agent-layer-for-youtube-advertising-and-creator-intelligence Accessed 2026-11-10.

[978] Your Brand Suitability Solution Claims to be Based on AI, But is it?. Pixability, Inc. (Kelly Patterson), 2023-05-25. https://www.pixability.com/resources/blog/your-brand-suitability-solution-claims-to-be-based-on-ai-but-is-it Accessed 2026-11-10.

[979] Pixability Announces Release of Major YouTube Brand Safety and Suitability Study in Collaboration with Global Alliance for Responsible Media. Pixability, Inc., 2021-10-12. https://www.pixability.com/about/news/pixability-announces-release-of-major-youtube-brand-safety-and-suitability-study-in-collaboration-with-global-alliance-for-responsible-media Accessed 2026-11-10.

[980] Introducing Pixie by Pixability (Pixie agentic AI). Pixability, Inc., undated. https://www.pixability.com/youtube-solutions/tech/pixie-ai Accessed 2026-11-10.

[981] Pixability Announces Wide Release Of BrandShield, AI-Powered Brand Suitability Solution For YouTube & YouTube On TV. Pixability (via PR Newswire), 2020-10-14. https://www.prnewswire.com/news-releases/pixability-announces-wide-release-of-brandshield-ai-powered-brand-suitability-solution-for-youtube--youtube-on-tv-301152260.html Accessed 2026-11-10.

[982] Pixability Updates Its Contextual Tools To Analyze YouTube With Generative AI. AdExchanger (Victoria McNally), 2024-10-07. https://www.adexchanger.com/ai/pixability-updates-its-contextual-tools-to-analyze-youtube-with-generative-ai/ Accessed 2026-11-10.

[983] Exclusion Lists: Avoiding the Worst of the Worst & Most Likely to Serve. Pixability, Inc. (Zoe Zimman), 2025-08-06. https://www.pixability.com/resources/blog/exclusion-lists-avoiding-the-worst-of-the-worst-and-most-likely-to-serve Accessed 2026-11-10.

[984] Pixability Customer Terms and Conditions (archived copy). Pixability, Inc., undated (no effective date shown) (event captured 2026-05-17 17:54:53 UTC). https://web.archive.org/web/20260517175453/https://www.pixability.com/terms-conditions/ Accessed 2026-11-10.

[985] Pixability Unveils 100% Brand-Safe Spend Guarantee for Video Advertising on YouTube. Pixability, Inc., 2018-02-07. https://www.pixability.com/about/news/pixability-unveils-100-brand-safe-spend-guarantee-for-video-advertising-on-youtube Accessed 2026-11-10.

[986] Pixability Launches New Contextual Solution to Solve for Transparency and Brand Suitability on Connected TV. Pixability, Inc., 2022-11-16. https://www.pixability.com/about/news/pixability-launches-new-contextual-solution-to-solve-for-transparency-and-brand-suitability-on-connected-tv Accessed 2026-11-10.

[987] Pixability Brings CTV Brand Suitability to the GroupM Premium Marketplace with new GroupM and IRIS.TV Partnership. Pixability, Inc., 2023-08-03. https://www.pixability.com/resources/blog/pixability-brings-ctv-brand-suitability-to-the-groupm-premium-marketplace-with-new-groupm-and-iristv-partnership Accessed 2026-11-10.

[988] Innovid & IRIS.TV Partner to Enable Actionable, Video-Level Contextual Measurement for CTV. IRIS.TV, 2024-12-18. https://blog.iris.tv/en/innovid-iris.tv-partner-to-enable-actionable-video-level-contextual-measurement-for-ctv Accessed 2026-11-10.

[989] IRIS_ID Data API for SSPs (IRIS.TV Knowledge Base). IRIS.TV, undated (page footer: Copyright 2026, Viant Technology). https://support.iris.tv/iris_id-data-api-for-ssps Accessed 2026-11-10.

[990] YouTube Activation Partner Roster (Google Enterprise Marketing Portal; Google-served script bundle). Google, live bundle. https://www.gstatic.com/alkali/717fcb439962164fd735614c0417d2ea3b119a30.js Accessed 2026-11-10.

[991] YouTube Measurement Program (archived copy, 2 August 2026). YouTube (Google), undated (captured 2026-08-02) (event 2026-08-02). https://web.archive.org/web/20260802003800/https://www.youtube.com/measurementprogram/ Accessed 2026-11-10.

[992] Pixability helps Travel Texas exceed reach and engagement goals across multiple ad formats (success story). Pixability, Inc., undated (sitemap lastmod 2025-09-29). https://www.pixability.com/resources/success-stories/pixability-helps-travel-texas-exceed-reach-and-engagement-goals-across-multiple-ad-formats Accessed 2026-11-10.

[993] Website Terms of Use. Pixability, Inc., 2026-10. https://www.pixability.com/terms-conditions Accessed 2026-11-10.

[994] Canela Media and Pixability Partner to Expand Hispanic Audience Reach on YouTube. Canela Media and Pixability, Inc., 2025-10-22. https://www.pixability.com/about/news/canela-media-and-pixability-partner-to-expand-hispanic-audience-reach-on-youtube Accessed 2026-11-10.

[995] Our Partners (Partner Connect Program). Pixability, Inc., undated. https://www.pixability.com/about/partners Accessed 2026-11-10.

[996] Inclusive Media Initiative for Creators. Pixability, Inc., undated. https://www.pixability.com/responsible-media/inclusive-media-creators Accessed 2026-11-10.

[997] Pixability Launches New Solution to Help Advertisers Avoid 'AI Slop' on YouTube. Pixability, Inc., 2026-10-01. https://www.pixability.com/about/news/pixability-launches-new-solution-to-help-advertisers-avoid-ai-slop-on-youtube Accessed 2026-11-10.

[998] Pixability and Ad Fontes Media Join Forces to Empower Advertisers to Safely Support News Content on YouTube. Pixability, Inc., 2024-09-25. https://www.pixability.com/about/news/pixability-and-ad-fontes-media-join-forces-to-empower-advertisers-to-safely-support-news-content-on-youtube Accessed 2026-11-10.

[999] Web Crawler (Comscore crawler information page). Comscore, Inc., undated. https://www.comscore.com/Web-Crawler Accessed 2026-11-10.

[1000] Keyword Targeting FAQs (Comscore Activation help centre). Comscore, Inc. (Proximic by Comscore), 2025-02-24 (event updated 2025-12-17 (API metadata)). https://comscoreactivation.zendesk.com/hc/en-us/articles/34340969436699-Keyword-Targeting-FAQs Accessed 2026-11-10.

[1001] Brand Safety and Suitability FAQs (Comscore Activation help centre). Comscore, Inc. (Proximic by Comscore), 2025-02-24 (event created 2025-02-24T22:58:50Z; updated 2025-02-24T23:00:32Z (API metadata)). https://comscoreactivation.zendesk.com/hc/en-us/articles/34341020054043-Brand-Safety-and-Suitability-FAQs Accessed 2026-11-10.

[1002] Comscore and MediaMath Partner to Launch Connected TV and Livestreaming Programmatic Contextual Targeting (press release). Comscore, Inc., 2020-12-03. https://www.comscore.com/layout/set/proximic/Insights/Press-Releases/2020/12/Comscore-and-MediaMath-Partner-to-Launch-Connected-TV-and-Livestreaming-Programmatic-Contextual-Targeting Accessed 2026-11-10.

[1003] Audio, Activated: Expanding the Intelligence Layer for Omnichannel Advertising (blog). Comscore, Inc. (Proximic by Comscore), 2026-08-04. https://www.comscore.com/layout/set/proximic/Insights/Blog/Audio-Activated Accessed 2026-11-10.

[1004] Contextual Targeting Solution (Comscore Content Activation). Comscore, Inc., undated. https://www.comscore.com/Products/Activation/Contextual-Targeting-Solution Accessed 2026-11-10.

[1005] Brand Protection: Achieve optimal campaign protection without sacrificing reach and ROI. Comscore, Inc., undated. https://www.comscore.com/Products/Activation/Contextual-Targeting-Solution/Brand-Protection Accessed 2026-11-10.

[1006] Inconsistencies in Classification of Online News Articles: A Call for Common Standards in Brand Safety Services (arXiv abstract page). Michael Smith, Riley Grossman, Antonio Torres-Aguero, Pritam Sen, Cristian Borcea, Yi Chen, 2026-01-03. https://arxiv.org/abs/2601.01303v1 Accessed 2026-11-10. *Access: abstract only.*

[1007] Epsilon Partners with Comscore's Proximic for Global Contextual Data (press release). Comscore, Inc. (Proximic by Comscore) and Epsilon, 2024-03-06 (event the Resources listing dates it 12 Mar 2024; the dateline reads March 6, 2024). https://www.comscore.com/layout/set/proximic/Insights/Press-Releases/2024/3/Epsilon-Partners-with-Comscores-Proximic-for-Global-Contextual-Data Accessed 2026-11-10.

[1008] Yahoo and Comscore Partner to Bring Advanced Brand Protection, Contextual Solutions to CTV Campaigns (press release). Yahoo Inc., 2022-08-18. https://www.yahooinc.com/press/yahoo-and-comscore-partner-to-bring-advanced-brand-protection-contextual-solutions-to-ctv-campaigns Accessed 2026-11-10.

[1009] Terms of Service (comscore.com). Comscore, Inc., 2021-02-09 (event 'Last modified: 09 February 2021'). https://www.comscore.com/About/Terms-of-Service Accessed 2026-11-10.

[1010] Audio at Parity: Comscore Launches Transcript-Level Targeting and Measurement with Spotify, SiriusXM, Triton Digital, Acast, and Libysn (press release). Comscore, Inc., 2026-07-22. https://www.comscore.com/layout/set/proximic/Insights/Press-Releases/2026/7/Audio-at-Parity-Comscore-Launches-Transcript-Level-Targeting-and-Measurement-with-Spotify-SiriusXM-Triton-Digital-Acast-and-Libysn Accessed 2026-11-10.

[1011] Comscore Announces Partnership with Amazon DSP to Expand Content Addressability to Drive Campaign Performance (press release). Comscore, Inc., 2026-06-24. https://www.comscore.com/layout/set/proximic/Insights/Press-Releases/2026/6/Comscore-Announces-Partnership-with-Amazon-DSP Accessed 2026-11-10.

[1012] MRC Announces Changes to Accreditation Statuses of Comscore Digital Measurement Products (statement). Media Rating Council (MRC), 2022-04-26. https://www.mediaratingcouncil.org/sites/default/files/News/Accreditation-Matters/042622%20MMX%20Hiatus%20Statement_Final.pdf Accessed 2026-11-10.

[1013] Privacy Centric-Targeting - Proximic by Comscore (home page). Comscore, Inc. (Proximic by Comscore), undated (panel figures stamped 'Data as of Feb. 2026'). https://www.proximic.com/home Accessed 2026-11-10.

[1014] Comscore Launches Audio Targeting and Measurement Capabilities with The Trade Desk (press release). Comscore, Inc., 2026-01-08. https://www.comscore.com/layout/set/proximic/Insights/Press-Releases/2026/1/Comscore-Launches-Audio-Targeting-and-Measurement-Capabilities-with-The-Trade-Desk Accessed 2026-11-10.

[1015] comScore, Inc. FY2025 10-K, R12: Revenue Recognition (disaggregation by solution group). comScore, Inc. (SEC filing, XBRL rendering), 2026-03-26. https://www.sec.gov/Archives/edgar/data/1158172/000115817226000009/R12.htm Accessed 2026-11-10.

[1016] Comscore Introduces Global Verified Contextual Private Marketplaces (PMPs) for Connected TV and Video Advertising (press release). Comscore, Inc., 2020-10-22. https://www.comscore.com/layout/set/proximic/Insights/Press-Releases/2020/10/Global-Verified-Contextual-Private-Marketplaces-for-Connected-TV-and-Video-Advertising Accessed 2026-11-10.

[1017] Horizon Media Taps Comscore as Preferred Contextual Intelligence Provider to Tackle Signal Loss and Privacy Challenges (press release). Comscore, Inc. (Proximic by Comscore) and Horizon Media, 2024-08-14. https://www.comscore.com/layout/set/proximic/Insights/Press-Releases/2024/8/Horizon-Media-Taps-Comscore-as-Preferred-Contextual-Intelligence-Provider-to-Tackle-Signal-Loss-and-Privacy-Challenges Accessed 2026-11-10.

[1018] comScore, Inc. FY2025 10-K, R50: Revenue Recognition - Schedule of Disaggregation of Revenue (geography, timing). comScore, Inc. (SEC filing, XBRL rendering), 2026-03-26. https://www.sec.gov/Archives/edgar/data/1158172/000115817226000009/R50.htm Accessed 2026-11-10.

[1019] Accredited Services: Television. Media Rating Council (MRC), undated live page. https://mediaratingcouncil.org/accreditation/television Accessed 2026-11-10.

[1020] For Media Sellers - Proximic by Comscore. Comscore, Inc. (Proximic by Comscore), undated (footer: © 2026 Comscore, Inc.). https://www.comscore.com/home/Programmatic-Targeting-Solutions/For-Media-Sellers Accessed 2026-11-10.

[1021] comScore, Inc. Definitive Proxy Statement (DEF 14A) for the 16 June 2026 annual meeting. comScore, Inc. (SEC filing), 2026-04-30 (event accepted 2026-04-30 16:01:05 (filing index)). https://www.sec.gov/Archives/edgar/data/1158172/000119312526197109/d82773ddef14a.htm Accessed 2026-11-10. *Access: partial extract.*

[1022] General FAQs (Comscore Activation help centre). Comscore, Inc. (Proximic by Comscore), 2025-02-24 (event updated 2025-12-17 (API metadata)). https://comscoreactivation.zendesk.com/hc/en-us/articles/34340824102939-General-FAQs Accessed 2026-11-10.

[1023] Proximic Activation Platform Release Notes (Comscore Activation help centre). Comscore, Inc. (Proximic by Comscore), 2026-04-27. https://comscoreactivation.zendesk.com/hc/en-us/articles/49667704508955-Proximic-Activation-Platform-Release-Notes Accessed 2026-11-10.

[1024] Sightly announces SmartLists 2.0 on its public company page. Sightly Enterprises, Inc., 2026. https://www.linkedin.com/posts/sightly_say-goodbye-to-static-lists-smartlists-activity-7462134339284189184-Ei-5 Accessed 2026-11-11.

[1025] Sightly SmartLists 2.0 release reproduced by Adtech Juice. Sightly; reproduced by Adtech Juice / The Digital Voice PR Agency Ltd, 2026-05-13 (event 2026-05-12). https://www.adtechjuice.com/sightly-launches-smartlists-2-0 Accessed 2026-11-11.

[1026] Solutions: Choose Your Path. Sightly Enterprises, Inc., undated. https://www.sightly.com/solutions/ Accessed 2026-11-10.

[1027] Sightly and Horizon Media Score Over 99% in IAS YouTube Brand Safety Measurement Beta. Sightly (press release via PRWeb), 2019-01-14. https://www.prweb.com/releases/sightly-and-horizon-media-score-over-99-in-ias-youtube-brand-safety-measurement-beta-887602555.html Accessed 2026-11-10.

[1028] MAGNA-Sightly Joint Study Reveals Authentic Cultural Alignment and Quick Action is Key to Increasing Branding KPIs. Sightly Enterprises, Inc. (study landing page), undated. https://www.sightly.com/magna-sightly-joint-study-reveals-authentic-cultural-alignment-and-quick-action-is-key-to-increasing-branding-kpis/ Accessed 2026-11-10. *Access: partial extract.*

[1029] Solving CTV's Blindspot: How Sightly + IRIS.TV Deliver Contextual Accuracy & Precision at Scale. Sightly Enterprises, Inc. (blog, by Kevin Barnes), 2025-04-30. https://www.sightly.com/solving-ctvs-blindspot-how-sightly-iris-tv-deliver-contextual-precision-at-scale/ Accessed 2026-11-10.

[1030] Sightly Brings Video-Level Curation to CTV, Ending the 'Black Box' of Streaming Ad Placement. Sightly Enterprises, Inc. (press release), 2026-08-12. https://www.sightly.com/news/sightly-brings-video-level-curation-to-ctv-ending-the-black-box-of-streaming-ad-placement/ Accessed 2026-11-10.

[1031] US11700218B2: Updating compatible distributed data files across multiple data streams of an electronic messaging service associated with various networked computing devices. Sightly Enterprises, Inc. (assignee); inventors Adam Eric Katz, Aman Raghuvanshi, Adam Jarrell Smith, Jacob Maximillian Miesner (Google Patents record of the USPTO grant), 2023-07-11 (grant) (event filed 2021-05-07). https://patents.google.com/patent/US11700218B2/en Accessed 2026-11-10.

[1032] Sightly Integrates with Google Display & Video 360, Bringing Brand Mentality Targeting to Buyers. Sightly Enterprises, Inc. (press release), 2026-09-30 (posting date; dateline 'New York, Sept 10') (event 2026-09-10 (announced availability)). https://www.sightly.com/news/sightly-integrates-with-google-display-video-360-bringing-brand-mentality-targeting-to-buyers/ Accessed 2026-11-10.

[1033] Terms of Use. Sightly Enterprises, Inc., 2025-04-04 (page states 'Updated 4th of April 2025'). https://www.sightly.com/terms-of-use/ Accessed 2026-11-10.

[1034] Sightly Segment Library. Sightly Enterprises, Inc., undated. https://www.sightly.com/segment-library/ Accessed 2026-11-10.

[1035] Brand Mentality - TikTok App Center. TikTok (listing); Sightly (app), listing created 2024-07-29 (createdAt 1722289112), updated 2024-10-01 (updatedAt 1727808393). https://ads.tiktok.com/app-center/detail/Brand-Mentality/7395712238868398097 Accessed 2026-11-10.

[1036] Sightly Expands YouTube Intelligence with Creator and Content Insights API Integration. Sightly Enterprises, Inc. (press release), 2026-10-29. https://www.sightly.com/news/sightly-expands-youtube-intelligence-with-creator-and-content-insights-api-integration/ Accessed 2026-11-10.

[1037] Sightly Integrates with IRIS.TV to Change the CTV Buying Game. Sightly (press release, EZ Newswire), 2023-08-15. https://www.sightly.com/news/sightly-integrates-with-iris-tv-to-change-the-ctv-buying-game/ Accessed 2026-11-10.

[1038] How Moen Drove 117% Traffic Growth with Sightly's Self-Service YouTube Offering (case study). Sightly Enterprises, Inc., undated (case-study sitemap last modified 2026-06-17). https://www.sightly.com/case-studies/how-moen-drove-117-traffic-growth-with-sightlys-self-service-youtube-offering/ Accessed 2026-11-10.

[1039] United Airlines Identifies Real-Time Ad Moments with Sightly's Brand Mentality. Sightly Enterprises, Inc. (news post relaying MediaPost), 2025-08-26. https://www.sightly.com/news/united-airlines-identifies-real-time-ad-moments-with-sightlys-brand-mentality/ Accessed 2026-11-10.

[1040] Managed Service Media. Sightly Enterprises, Inc., undated. https://www.sightly.com/managed-service-media/ Accessed 2026-11-10.

[1041] AI Digital Invests in and Partners with Sightly to Infuse High-Value Programmatic Campaigns with Brand Mentality. Sightly Enterprises, Inc. (repost of AI Digital release), 2023-07-18. https://www.sightly.com/news/ai-digital-invests-in-and-partners-with-sightly-to-infuse-high-value-programmatic-campaigns-with-brand-mentality/ Accessed 2026-11-10.

[1042] Form D: Notice of Exempt Offering of Securities, Sightly Enterprises, Inc. (accession 0000905718-21-001280). Sightly Enterprises, Inc. (filer); U.S. Securities and Exchange Commission (EDGAR), 2021-10-06 (event first sale 2018-12-31). https://www.sec.gov/Archives/edgar/data/1590014/000090571821001280/primary_doc.xml Accessed 2026-11-10.

[1043] Data Protection & Security Reference. Sightly Enterprises, Inc., undated. https://www.sightly.com/data-protection-security/ Accessed 2026-11-10.

[1044] Is YouTube billing - and not refunding - advertisers for ads served on non-conforming or potentially copyright infringing content? (Executive Summary PDF). Adalytics Research LLC, 2025-07. https://drive.proton.me/urls/X38XT9AN34#56PenOxcKUpX Accessed 2026-11-10.

[1045] TikTok partners with Zefr on brand safety and brand suitability measurement (TikTok For Business blog). TikTok, 2022-01-24. https://ads.tiktok.com/business/en-US/blog/zefr-partnership-brand-safety-suitability-measurement Accessed 2026-11-10.

[1046] Brand Safety (Reddit Ads Help). Reddit, Inc.. https://business.reddithelp.com/s/article/Reddit-Ads-Brand-Safety Accessed 2026-11-10.

[1047] Taboola strengthens its position on Brand Safety with Adloox. Adloox, 2023-02-01. https://www.adloox.com/blog/taboola-strengthens-its-position-on-brand-safety-with-adloox Accessed 2026-11-10.

[1048] Renewed Media Rating Council Accreditations for Adloox. Adloox, 2023-05-15. https://www.adloox.com/blog/media-rating-council-accreditations-for-adloox-renewed-again Accessed 2026-11-10.

[1049] White Ops is the First Company to Receive MRC Accreditation for Both Pre-Bid and Post-Bid Invalid Traffic Detection and Mitigation Across Desktop, Mobile Web, Mobile In-App and Connected TV (press release), Internet Archive capture of 20 October 2021. White Ops (now HUMAN), 2021-01-27. https://web.archive.org/web/20211020005029id_/https://www.humansecurity.com/newsroom/white-ops-is-the-first-company-to-receive-mrc-accreditation-for-both-pre-bid-and-post-bid-invalid-traffic-detection-and-mitigation-across-desktop-mobile Accessed 2026-11-10.

[1050] Vista Equity Partners page ('Ushering in a New Age of Ad Innovation'). Integral Ad Science, 2025-06-03. https://integralads.com/about-ias/vista-equity-partners/ Accessed 2026-11-10.

[1051] IAS earns industry's first MRC accreditation for CTV Viewable Impressions (IAS Insider). Integral Ad Science, 2023-05-18. https://integralads.com/insider/ias-first-media-rating-council-accreditation-ctv-viewable-impressions/ Accessed 2026-11-10.

[1052] IAS Earns MRC Accreditation for Third-Party Calculation and Reporting of YouTube Viewability (IAS Insider). Integral Ad Science, 2024-03-21. https://integralads.com/insider/mrc-accreditation-youtube-viewability/ Accessed 2026-11-10.

[1053] Integral Ad Science Receives Industry's First Ethical AI Certification from Alliance for Audited Media (release). Integral Ad Science, 2025-07-29. https://integralads.com/news/ias-receives-industrys-first-ethical

-ai-certification-from-alliance-for-audited-media/ Accessed 2026-11-10.

[1054] Integral Ad Science Appoints Lidiane Jones Chief Executive Officer (release). Integral Ad Science, 2026-07-07. https://integralads.com/news/integral-ad-science-appoints-lidiane-jones-chief-executive-officer/ Accessed 2026-11-10.

[1055] IAS Reports Third Quarter 2025 Financial Results (Exhibit 99.1 to Form 8-K, Item 2.02, filed November 4, 2025). Integral Ad Science Holding Corp. (SEC filing exhibit), 2025-11-04. https://www.sec.gov/Archives/edgar/data/1842718/000184271825000107/q325earningsrelease.htm Accessed 2026-11-10.

[1056] Pixalate Receives MRC Accreditation (display ad viewability). Pixalate, Inc., 2014-11-25. https://www.pixalate.com/blog/mrc-accreditation-announcement Accessed 2026-11-10.

[1057] Pixalate Receives MRC Accreditation for Sophisticated Invalid Traffic (SIVT) Detection. Pixalate, Inc., 2017-07-05. https://www.pixalate.com/blog/mrc-accreditation-sivt-invalid-traffic-givt-detection-filtration-mobile Accessed 2026-11-10.

[1058] Pixalate adds 12 new MRC accredited metric areas; becomes first fraud detection company accredited for SIVT filtration across display, in-app, video and OTT. Pixalate, Inc., 2018-11-28. https://www.pixalate.com/blog/12-new-mrc-accreditations-desktop-mobile-app-ott-ad-fraud Accessed 2026-11-10.

[1059] Pixalate receives continuation of MRC accreditation for SIVT detection and filtration across digital environments. Pixalate, Inc., 2020-05-13. https://www.pixalate.com/blog/mrc-accreditation-multiple-environments-sivt Accessed 2026-11-10.

[1060] LaMonica v. Pixalate, Inc., Adv. No. 21-01018 (Bankr. S.D.N.Y.), docket. CourtListener (RECAP), 2021-02-09 (event 2021-02-09 to 2022-05-09). https://www.courtlistener.com/docket/59240396/lamonica-v-pixalate-inc/ Accessed 2026-11-10. *Access: partial extract.*

[1061] Pixalate renews all 23 MRC accreditation areas. Pixalate, Inc., 2023-02-16. https://www.pixalate.com/blog/pixalate-renews-all-23-mrc-accreditation-areas Accessed 2026-11-10.

[1062] Form D notice of exempt offering, ZEFR, INC. (accession 0001474764-13-000002). ZEFR, INC. (filed with the SEC), 2013-01-16 (event 2012-12-18). https://www.sec.gov/Archives/edgar/data/1474764/000147476413000002/xslFormDX01/primary_doc.xml Accessed 2026-11-10.

[1063] Zefr Selected By Meta as First Partner for Feed Verification. Zefr (press page), 2022-03-16. https://zefr.com/press/zefr-selected-by-meta-as-first-partner-for-feed-verification Accessed 2026-11-10.

[1064] GARM expands membership to independent brand safety providers (repost). Zefr (press page), 2022-06-16. https://zefr.com/press/garm-expands-membership-to-independent-brand-safety-providers Accessed 2026-11-10.

[1065] Zefr Launches Atrium, A New Measurement Product Suite (PR Newswire, November 2022). Zefr (press page), 2022-11-03. https://zefr.com/press/zefr-launches-atrium-a-new-measurement-product-suite-providing-advertisers-with-fully-transparent-brand-safety-and-suitability-measurement-across-walled-gardens Accessed 2026-11-10.

[1066] Zefr + Meta Brand Suitability Verification now Available for Facebook Feed, Powered by AI (PR Newswire, March 2023). Zefr (press page), 2023-03-30. https://zefr.com/press/zefr-meta-brand-suitability-verification-now-available-for-facebook-feed-powered-by-ai Accessed 2026-11-10.

[1067] Zefr + TikTok Pilot Video Exclusion Controls and Expand Brand Suitability Placement Availability (PR Newswire, October 2024). Zefr (press page), 2024-10-08. https://zefr.com/press/zefr-tiktok-pilot-video-exclusion-controls-and-expand-brand-suitability-placement-availability Accessed 2026-11-10.

[1068] Zefr Expands Pre-Impression Content Avoidance for Meta Feed and Reels, Now Globally Available (PR Newswire, February 2025). Zefr (press page), 2025-02-13. https://zefr.com/press/zefr-expands-pre-impression-content-avoidance-for-meta-feed-and-reels-now-globally-available Accessed 2026-11-10.

[1069] Zefr Expands Strategic Partner Network Across Türkiye, GCC and Egypt with Black C Media and RedC Media. Zefr (press page), 2026-07-13. https://zefr.com/press/zefr-expands-strategic-partner-network-across-t%C3%BCrkiye-gcc-and-egypt Accessed 2026-11-10.

Verified Against What — a research paper by Evgeny Popov, No Fluff Advisory. Research cutoff 11 November 2026.

Built from 10 research workstreams, a vendor-universe sweep and 7 company profiles, each checked by a separate AI verifier instructed to refute it, after an outside AI review of the brief whose every factual point was checked at its primary source. AI-assisted research, one scoring pass, not externally audited or peer reviewed.