Media Rating Council, Inc.

420 Lexington Avenue
Suite 343
New York, NY 10170

Tel: (212) 972-0300
Fax: (212) 972-2786
www.mediaratingcouncil.org

April 11, 2025

MRC IVT Requirements and Processes for Pre-Bid – Comments on the Recent Adalytics Blog

A recent blog post by an ad tech vendor, Adalytics, shared observations related to ad serving activity to robotic agents (bots) and observations related to certain Media Rating Council (MRC) accredited vendors (specifically, DoubleVerify/DV, HUMAN and Integral Ad Science/IAS) as well as made statements about MRC invalid traffic (IVT) requirements and processes. The blog was covered in various media publications.

It should be noted that MRC codifies requirements related to invalidity/validity and does not use the term fraud, as fraud is a legal term that implies intent that cannot always be discerned and is a narrower definition than validity. Often times, invalid traffic occurs without fraudulent intent (such as crawling or scraping content for classification), but this traffic is invalid for ad measurement. The bots mentioned in the blog post likely do not have fraudulent intent to generate ad revenue, but exist for benign purposes such as archiving web pages.

The MRC's review of this blog post found omissions related to and an incomplete understanding of MRC's IVT requirements as well as incomplete information related to audited vendor measurement and processes that may lead readers to incorrect conclusions. While MRC does not frequently comment publicly on audit matters, as our process involves confidential information about audited vendors (fully reviewed by our members), we believe the incomplete presentation of details related to these matters required us to comment, without compromising confidentiality.

MRC became aware of Adalytics' potential blog post as early as December 2024 and throughout early 2025, based on contacts for comment on it by various publications who were apparently informed about the content of the blog (at no point was the blog shared with MRC before publication). MRC did not offer comment with this type of severely limited context, however we did share information on our general requirements as background.

In summary, we had the following observations related to the blog post, all of which are explained further later in this memorandum:

-MRC requires back end (post-serve) filtration of General Invalid Traffic (GIVT) on an impression by impression basis for all accredited digital vendors; this includes known Spiders and Bots as well as known dedicated Data Center IP addresses.

-MRC does not require pre-bid (upfront or pre-serve) filtration or blocking of IVT, in fact, our Standards discourage or caution against it in several instances. Full back end filtration is required to inform pre-bid and must be applied to all measured and reported impressions, regardless of pre-bid usage.

-MRC audits (across multiple periods over the course of recurring annual audits) for accredited digital vendors include testing of the application of the IAB Spiders and Bots list, (we noted that the declared bot referenced in the blog post is included in the applied IAB Spiders and Bots list), as part of required back end GIVT processes in production as applied to measured and reported impressions.

-MRC audits (across multiple periods over the course of recurring annual audits) for accredited digital vendors include testing of the application of industry data center lists such as the TAG Data Center IP list, including the known dedicated data center IP addresses, as part of required back end GIVT processes in production as applied to measured and reported impressions.

-Where pre-bid IVT detection is accredited, MRC audits (across multiple periods over the course of recurring annual audits) include testing of application of back end observations, including the IAB Spiders and Bots list and industry data center lists (e.g., the TAG Data Center IP list) used to inform bid request classification process and pre-bid files and databases.

-None of this is either based on sampling or reliant on DSPs or SSPs sharing or "passing" information (such as IP Address, User Agent or Device ID) to the vendors to create pre-bid assets for bid consideration.

-The decision to present a bid request or serve an ad solely rests with DSPs and SSPs and is dependent on what information they choose to compare to or assess within pre-bid products.

-HUMAN provides suggestions on whether a request is IVT or not based on back-end detection and filtration through an API. This is only provided to DSPs and SSPs, not advertisers or agencies. HUMAN does not directly filter requests, prevent ad serving or enable blocking, and the decisions to bid or not bid on these requests based on HUMAN's suggestions are solely made at the DSP/SSP level and outside of our audits.

-DV provides pre-bid segments and IVT classifications to DSPs, SSPs and advertisers/agencies through DSPs which may be queried through either an API configuration or a flat-file to help these organizations decide whether or not to filter a request for IVT or not based on back-end detection data, which is delivered through either a file upload or API. DV does not directly filter requests or prevent ad serving, but does enable blocking post-serve only where configured by DV's clients (which is evaluated as part of pre-bid audits).

It should be noted that Adalytics did not reach out to MRC to seek clarification on our requirements at any point during the compilation of their blog post, nor did they present perceived shortcomings of our process to us. We generally welcome such outreach, even from unaudited/unaccredited vendors, and if verifiable information leads to the need to improve or update our requirements or processes, we strive to reflect this in updated requirements as evidenced by the many updated Standards and Guidelines issued by MRC (including the IVT Standards which have been updated and supplemented with Interim Updates many times since their initial release to reflect new observations).

Adalytics is not audited or accredited by MRC so their processes of discovery and testing and, most importantly, their projection of the impact of findings (or implied impacts in most cases) have not been reviewed by MRC for controls and appropriateness.

MRC IVT Requirements

Related to this specific blog post, Adalytics references certain excerpts of MRC's IVT Standard (the first and only IVT Standard that has been produced in our industry to guide the detection and filtration of IVT) including aspects related to known spiders/bots and known invalid data center traffic. These are General Invalid Traffic (GIVT) requirements of all MRC audited and accredited digital measurement organizations, not just IVT focused organizations. It should be noted that GIVT is IVT that can be identified by list-based or parameter driven techniques and MRC requires multiple types of GIVT be detected beyond known bots and data centers, not all GIVT "self-declares" and not all GIVT is "benign" such as content scrapers, but it is all invalid for ad measurement.

The first salient point related to this is that the MRC requires **back end (post-serve) filtration of GIVT. The MRC does not require pre-bid (front-end) filtration of such activity, in fact our Standards discourage it in several instances.**

The MRC IVT Standards Sections 4.2 and 4.2.1, which were not referenced in Adalytics' blog post, directly cover front-end IVT techniques and our requirements of them.

https://mediaratingcouncil.org/sites/default/files/Standards/IVT%20Addendum%20Update%20062520.pdf

An excerpt from these sections states (our emphasis added):

*"**Up-front detection techniques where a bid request is not fulfilled or otherwise blocked due to IVT must be employed with caution because they are particularly prone to telegraphing detection techniques, in most cases, to the traffic source because of an element of blocking that becomes apparent.** As such, they tend to become less effective over time without additional research and development into new detection methodologies. **Application of these techniques on an up-front basis is not required. Back-end detection and removal techniques are more invisible to the source, and therefore less prone to signaling detection methodologies to IVT perpetrators, however they may add complexity to reporting and processing, since data that flows through the measurement organization will therefore contain IVT transactions that are removed in later stages.** See further guidance on up-front techniques in Section 4.2.1 below."*

While the MRC has reservations about the risks and effectiveness of pre-bid approaches to IVT, in our 2020 IVT Standards update, we sought to standardize requirements for use of these approaches and to enable an audit framework as these approaches are desired, often by ad serving organizations. Section 4.2.1 goes on to summarize these requirements. This includes the stipulation that back-end detection and filtration techniques are required for compliance with the Standard and that digital measurement organizations employing up-front IVT filtration techniques must do so in combination with required back-end detection and filtration techniques.

It should be noted that MRC recognizes the desire for pre-bid approaches in the industry for efficiency of budgets, assurance of demonetization, carbon considerations and

avoidance reasons. MRC also recognizes the thought that known GIVT pre-bid filtration may pose diminished risk of reverse engineering. However, IVT decisions are often made holistically and not distinguished between GIVT/SIVT. MRC holds that pre-bid filtration poses risks detailed above and may also mask the occurrence of IVT in back-end reporting, limiting the ability for measurement users to identify properties with higher levels of invalid traffic. Further, while a known bot in the GIVT category declaring or using a known data center may be widely known to be invalid, blocking (which often is based on a range of IVT decisions and telemetry including SIVT) may allow that bot to collect information about decisions or a vendors' detection techniques at scale to A/B test for IVT weaknesses. For this reason, MRC continues to discourage pre-bid filtration or blocking without care and does not require it. See further discussion below regarding potentially lower-risk DSP/SSP blocking.

The Standard also states that it is understood that a measurement organization may not have full visibility into how their up-front resources are deployed by an ad-server, platform or network and as a result, may not be able to directly quantify the impact on requests, bids, blocks or impressions. This is a key aspect that may not be fully explored in Adalytics' blog that we will discuss later.

Finally, the Standard states that in certain cases, measurement organizations may employ up-front techniques where a bid is pre-emptively invalidated based on IVT, but is still fulfilled with an impression served. Such approaches may obviate some of the requirements of up-front approaches discussed above such as those related to preventing reverse engineering (as the bid request would still be fulfilled and not known to a potential bad actor) and disclosure as filtration of impressions would still be reported. Again, a key distinction that does not appear to be fully considered in Adalytics' blog, but may actually account for a large majority of its observations.

If up-front techniques are based on list-based resources developed and deployed based on back-end techniques such as spiders and bots, they may be considered deterministic and require less support than probabilistic techniques.

MRC Accreditation of Vendors

MRC accreditation is based on annual audits by independent CPAs under AICPA's Attestation Standards and reviewed by our members that includes review and inspection of code, tests of controls, data analytics and designed activity testing in production environments. It also includes visibility into production data collection and reported results which cannot be gleaned from external processes.

DV, HUMAN and IAS are all MRC audited and accredited for **back end** detection and filtration of GIVT (required of all accredited digital services) and Sophisticated Invalid Traffic (SIVT; an encouraged but not required aspect of MRC's Standard that involves non-list based or non-parameter-driven techniques) for certain formats and environments specified on our site:

https://mediaratingcouncil.org/accreditation/digital

Back end IVT is expected to be applied for all traffic on an impression by impression basis. There is no sampling present in any of the IVT services MRC currently audits and accredits, but if sampling is applied, MRC also has methodological and error disclosure requirements.

In addition, DV and HUMAN are MRC audited and accredited for Pre-Bid IVT **detection (not prevention of ad serving or decisions to serve ads or not)** again for certain formats and environments detailed on our site.

IAS is not audited nor accredited for Pre-Bid IVT detection or their publisher tools mentioned in the blog post. IAS will be submitting it's Pre-Bid IVT detection to audit during 2025.

For DV, neither the Scibids targeting product nor the DV publisher tools mentioned in the blog post are audited or accredited.

Pre-Bid

Another issue we noted with Adalytics' blog post is the implication that accredited pre-bid IVT measurement is expected to result in direct prevention of ad serving and that a reader may conclude that observations of ads served to bots may be an indication that pre-bid detection is not functioning as represented or intended. Adalytics' blog recognizes this may not be the case as the Caveats and Limitations section at the end of the blog states:

"Thus, it can be difficult to determine or infer who was "responsible" or made a "decision" that resulted in an ad being served to a bot. Even if a vendor has correctly identified a user as a bot, sometimes other stakeholders may intentionally or inadvertently still authorize an ad to be served to that bot.

For example, a vendor may correctly identify a bot, and pass that information to another vendor who controls the ad serving or bidding decisioning process. That second vendor - for a number of reasons - may disregard, ignore, or fail to action the accurate and correct bot classification returned by the first vendor."

As recognized formally in our IVT Standard, often times, and as is the case for DV and HUMAN, a measurement vendor is not able to make ad serving decisions (to serve an ad or to respond to a bid in any way). Instead, both vendors provide assets for pre-bid detection that must be actioned by ad servers and configured by their clients.

Without compromising sensitive methodological details, in general, these services either provide files or databases that can be queried through API by ad serving organizations that list IP addresses, User Agents and other heuristics such as Device ID that have been observed to be invalid using back-end detection techniques (as required by MRC).

This functionality does not require "passing" of IP addresses or User Agents from ad serving organizations to the vendors to build pre-bid assets or lists, but it is based only on available or observable information using back end post-serve telemetry or observations (tagging at a site or ad level) and effectiveness of application of these assets depends on what a DSP or SSP chooses to compare the vendor assets to during decisioning. This is important to note as much of the reporting surrounding this blog post fails to understand that DSPs and SSPs generally communicate with each other, not ad tech vendors, and also continually mixes back end filtration with pre-bid detection. Open programmatic DSPs and SSPs generally utilize the IAB Tech Lab's OpenRTB protocol for communication

regarding requests and bids. User Agent and IP Address are part of this protocol, but recommended, not required fields, that may not be available or included for various reasons outside of MRC's purview.

Ad tech vendors are not reliant on DSPs or SSPs to collect things like User Agent or IP address for IVT consideration. Instead vendor tags are either present on properties (pages, apps, etc.) and/or are embedded within creatives delivered to them and are able to directly collect this information and apply it on the back end. However, once vendors produce pre-bid files/databases for use by DSPs or SSPs or are asked to evaluate a bid request based on these back end observations, the ad serving organization can choose to evaluate User Agent, IP Address or Device ID (all three, any combo of the three or none). The vendor's configuration and implementation guides allow and encourage utilizing all available information, but what a DSP or SSP chooses to use and compare is out of the vendor's control.

These files and databases are updated (near real time for database APIs and every several minutes for flat files) and made available to client users. For HUMAN, clients of their audited pre-bid product (Ad Fraud Defense) are generally ad tech platforms (DSPs and SSPs) and DV offers their audited pre-bid product to DSPs and SSPs, as well as to advertisers/agencies, which must be enabled and configured through their DSP.

When a DSP receives a bid request or when an SSP decides to send a bid request, they are able to cross-reference the pre-bid files or query the database APIs for an IVT determination based on back-end observations. It is ultimately the DSP's or SSP's decision (based on client configurations) on how to use and configure these pre-bid functions, what information they cross-reference and query and how to respond to or send a bid request and serve an ad. Neither DV nor HUMAN have control over these decisions. MRC's audit of vendor pre-bid services are limited to the compliant production, accuracy, functionality and availability of files and databases, and do not include upstream processes at DSPs or SSPs outside of the vendor's control or purview. MRC does not produce requirements for or have insight into billing or financial arrangements that may include billing for measurement (regardless of validity or delivery), ad serving or impressions (valid or invalid) per our mandate, and our audits focus on measurement and report of ad activity.

HUMAN does not enable post-serve blocking and their pre-bid metrics only classify requests as GIVT or SIVT. DV does enable post-serve blocking based on client configuration, but note this does not impact an ad server's decision to serve an ad, only restricts delivery of it. MRC's audits note a small % of ads are blocked as configured by DV clients using their pre-bid service, but delivered ads are monitored and subject to back end filtration. This functionality is evaluated as part of pre-bid audits where present.

It should also be noted that for commonly known GIVT, such as that listed within the IAB Spiders and Bots list or the TAG Data Center IP list, a DSP or SSP (and even publisher) can subscribe to and utilize these lists themselves in their decisioning without reliance on a vendor, and such use likely comes with minimal risk as it does not incorporate other IVT telemetry. However, vendors may provide more complete and comprehensive consideration of all IVT (all GIVT as well as SIVT). MRC also encourages third-party measurement of IVT where possible. MRC does not produce requirements for nor audit ad serving processes. We do audit certain organizations and services that may serve ads, but that is

limited to back-end measurement of delivery (and IVT filtration), not ad serving decisions nor pre-bid considerations.

Specific Bots

Adalytics' blog post referenced three distinct bots: HTTP Archive, an undisclosed bot and URLScan.io.

HTTP Archive is a declared bot present on the IAB Spiders and Bots list (since 2013) as including "PTST" in the User Agent string. MRC audits (regular recurring annual audits) test the requirement that accredited digital vendors apply the IAB Spiders and Bots list (updated regularly) in production for back end filtration including HTTP Archive as required by our IVT Standard. MRC audits also test the requirement that back end detection using this list inclusive of this bot is used to inform audited and accredited pre-bid products. MRC has not tested each of the example observations in Adalytics' blog post, but seeks to obtain reasonable assurance via observation and testing of application of the IAB bots list to measured and reported production traffic at the audited vendors. MRC believes the vendors likely have readily available information regarding each of the observations in the blog post and how they were classified with respect to IVT pre and post and encourages these vendors to share that information with clients and perhaps, publicly.

For the undisclosed bot, we cannot confirm the above as the identity of the bot was not shared in Adalytics' blog. If Adalytics would like to privately identify this to MRC, we can work to investigate its occurrence, however, Adalytics also notes that this bot does not declare its User Agent as robotic and is not listed on the IAB Spiders and Bots list. Adalytics' blog post mentions that this bot may use data centers (perhaps some dedicated known IVT data centers) and also may exhibit SIVT traits. More on that below.

Finally, the URLScan.io bot does not declare its User Agent as robotic and is not listed on the IAB Spiders and Bots list. However, the referrer Domain is declared for traffic generated by this bot. Our audits are able to inspect how traffic generated by this bot is treated by audited vendors in production as part of recurring audit processes. Adalytics' blog post mentions that this bot may use data centers (perhaps some dedicated known IVT data centers) and also may exhibit SIVT traits. More on that below.

Data Centers

Adalytics' blog post mentions that both the declared and undeclared bots often make use of well-known data center IPs. It's important to note that GIVT requirements and the associated TAG data center IP list used by accredited vendors, is limited to known invalid data center IPs dedicated to IVT or non-human activity, not mixed used data center IPs or unknown special purpose data center IPs that may require SIVT techniques to detect.

However, well-known cloud hosting data center IPs are included on TAG's list along with other known dedicated data center IPs.

MRC audits (regular recurring annual audits) test the requirement that accredited digital vendors apply industry data center lists (e.g., the TAG data center IP list; updated regularly) in production for back end filtration including known dedicated data centers as required by our IVT Standard.

MRC audits also test the requirement that back end detection using data center lists are used to inform audited and accredited pre-bid products. MRC has not tested each of the example observations in Adalytics' blog post, but seeks to obtain reasonable assurance via observation and testing of application of industry data center lists (e.g., the TAG data center IP list) to measured and reported production traffic at the audited vendors. MRC believes the vendors likely have readily available information regarding each of the observations in the blog post and how they were classified with respect to IVT pre and post and encourages these vendors to share that information with clients and perhaps, publicly.

SIVT

As mentioned above, Sophisticated Invalid Traffic, or SIVT, is not required for all digital measurement providers, but is strongly encouraged. SIVT has been audited and accredited for DV, IAS and HUMAN for backend detection and filtration and for DV and HUMAN for pre-bid detection for certain formats and environments specified on our site.

SIVT goes beyond IVT that can be detected and measured using list-based or parameter driven techniques, is far more difficult to detect and may vary from vendor to vendor. MRC does not require accredited SIVT vendors to identify all SIVT, as it's impossible to determine this with evolving sophistication of IVT techniques. Instead, MRC IVT Standards define multiple SIVT categories and techniques and require accredited SIVT vendors to assess risk in each category, and develop corresponding techniques and controls for each, assessed and tested for effectiveness by our audits. It's very likely a material amount of SIVT goes undetected and MRC pushes for continuous improvement through our processes and Standards to attempt to promote as thorough detection approaches as possible.

Pre-bid SIVT functions in much the same way as detailed above, with back end observations of SIVT associated with IPs, UAs and device IDs as well as properties used to inform pre-bid files and databases for DSP/SSP consideration in their serving decisions. Again, the vendors do not make these decisions, and MRC requires backend filtration and detection for SIVT across all categories even when pre-bid is present.

It's possible some of the observed activity in Adalytics' blog manifests as SIVT (recognized by Adalytics in the blog) if the User Agent is not properly declared as robotic, designed to spoof or simulate valid/human activity, mixed used or unknown data centers are used (or not at all) and through other means designed to mask this activity. While MRC thoroughly assesses accredited SIVT vendors across all required categories, we cannot definitively say all activity was properly detected as SIVT where applicable. MRC plans to incorporate some of the direct observations in our required testing as part of ongoing efforts, however, it should be noted that the focus of Adalytics' blog post was on missed basic GIVT and known activity.

MRC's Perspective on Adalytics:

MRC has no direct relationship with Adalytics; as previously noted they are not in the audit/accreditation process. Typically, MRC's only formal relationship with vendors is an audit and industry self-regulatory relationship but that does not exist with Adalytics. Further, MRC has no visibility into Adalytics' processes or services they provide to customers, the rigor of their research and the relevant standards compliance of their measurement. Given all this and the fact

that Adalytics positions itself as a vendor competitive to many vendors it often writes about, MRC takes care to consider their findings, as we would with any vendor.

While Adalytics' blog posts are often extensive with screenshots and detailed analysis that includes observations that can be confirmed, Adalytics also takes care at the end of each blog post to disclaim direct conclusions, they often state they cannot quantify impact of any findings or financial harm, intent or fault along with other stated caveats and limitations (an example of this for the blog subject to this release is noted above in the Pre-Bid section). However, they often pose questions or may lead a reader to imply issues that were not confirmed directly in their work or that MRC and our audits have found at times to be based on incomplete understanding of our requirements or processes. While Adalytics has published blog posts for some time, and much of these historical and even recent posts did not intersect with MRC audits as they involved unaudited/unaccredited services or aspects of measurement, over the past few years MRC has reviewed several that have been relevant to our activities. We actively review external research and unaudited vendor collateral for continuous self-assessment of our process and requirements and often times speak to many practitioners in our industry (audited or otherwise) to do the same.

MRC does believe that Adalytics' blogs promote discussion within the industry around important issues, focus on improvement and transparency and have been leveraged by MRC to seek audited services to improve rigor and transparency. We have also used these posts and other industry research to challenge our processes and requirements and see this impact as positive.

MRC has spent considerable time and effort considering the findings of Adalytics that intersect with our audits over time as shared with and reviewed by our members, and while we believe this has had some merit in terms of pushing for further rigor and transparency, we often times see little real impact from these findings related to our requirements or processes, unlike what Adalytics or other readers of their material sometimes imply or hypothesize.

MRC will continue to work to fulfill our mission in working across the industry to standardize measurement and independently audit services that voluntarily submit to our processes against these standards with the single goal of improving the validity, reliability and effectiveness of the measurement ecosystem. This includes transparency and education as it is clear that MRC can work harder to educate the industry regarding our requirements and processes and audited services can do more to educate the industry, beyond their customers, regarding how their services work including any limitations of measurement.

If you have any questions regarding this or would like to learn more about MRC and our specific requirements, our activities or membership in MRC, please contact us at staff@mediaratingcouncil.org.

About MRC

The Media Rating Council is a non-profit industry association established in 1963 comprised of leading television, radio, print and digital media companies, as well as advertisers, advertising agencies and trade associations, whose goal is to ensure measurement services that are valid, eliable and effective. Measurement services desiring MRC accreditation are required to disclose to their customers all methodological aspects of their service; comply with the MRC Minimum Standards for Media Rating Research as well as other applicable industry measurement guidelines; and submit to MRC-designed audits to authenticate and illuminate their procedures.

In addition, the MRC membership actively pursues research issues they consider priorities in an effort to improve the quality of research in the marketplace. Currently approximately 110 research products are audited by the MRC. Additional information about MRC can be found at www.mediaratingcouncil.org.

Media Contact
Bill Daddi, DBC Brand Communications
917-620-3717
bill@daddibrand.com

#
#
#