IAB Podcast Measurement Guidelines

Version 2.0

Draft for public comment

Please email Amit Shetty at podcastTech@iabtechlab.com with your feedback by 09/08/2017 (public comment period extended)

Drafted by the IAB Podcast Technical Working Group

Table of Contents

1. Executive Summary

Podcast audiences represent a growing segment of effective marketable media. Podcast ad revenues are expected to reach over $220 Million in 2017, an 85% increase from 2016, according to the IAB [Podcast Advertising Revenue Study](#) conducted by PwC US. Podcasts offer advertisers a hyper focused audience that chooses to listen to that content. Podcast listeners have been shown to be the most loyal and engaged audience of any digital medium. However, measurement of content and ad consumption in podcasting has not been consistent so far, which sometimes limits participation of brand advertisers. This document provides an introduction to tracking ad delivery in a podcast and attempts to provide clarity in the marketplace by describing best practices for measuring downloads, audience size, and ad delivery.

Podcasts are downloaded to a device for later listening or for online listening. In most cases the podcast file and any ads included with it are downloaded to a device that doesn't, or can't, send data about the consumption of the podcast and ads. This lack of data beyond ad delivery limits real-time measurement. In contrast, other media are consumed by reading an article and interacting with a site, playing a game, or streaming a video, all of which can be measured in real time. Even audio stations that offer music or news are streamed and measured in real time in today's media marketplace.

Podcast listeners have the ability to download files to consume whenever and wherever they want and are not required to have an active internet connection to play back an episode. The medium, the distribution, and the platforms used to collect and listen are built around the habit of downloading the file. Tracking content in this time-shifted medium involves filtering server logs to produce meaningful data for measurement. Since podcast technical teams analyze server logs differently, results vary across the industry.

The challenge for podcast producers and distributors is to offer buyers a set of metrics that is consistently defined and measured equally across the podcast medium. The definitions in this document aim to reduce measurement discrepancies and present a set of recommended metrics and guidelines based on industry best practices. With a consistent set of podcast advertising metrics, buyers and sellers can engage in a conversation about campaign strategy with confidence.

Audience

While all professionals in the podcast supply chain can benefit by being familiar with this document, metric definitions are primarily intended for podcast producers and distributors. Specifically, account managers should be familiar with and use metrics as defined in this document when negotiating ad packages with buyers. Additionally, podcast ad operations teams should use the metric definitions in this document to design or adjust the ad measurement technology they use to analyze server logs for podcast ad measurement.

Buyers should also reference this document to better understand how ads in podcast content are counted. This document offers a set of metrics that establish a mutual understanding in podcast advertising negotiations.

About this Version:

This document was originally released in September 2016. The updates in this version (v2.0) focus on improving the definitions for content metrics and providing a recommended process that helps improve the accuracy/correctness of the metrics.

2. Overview

Podcast content is an on-demand media format that listeners either download to listen to later or consume online. Unlike the streaming format more common in video, podcasts continue to be downloaded because of the convenience offered by existing platform and application functionality.

Despite the use of the word "streaming" in podcasting, "streamed" podcast files are progressively downloaded via the standard HTTP protocol. True streaming—typically reserved for live events—requires a specialized server and uses an entirely different protocol.

While "streaming" a podcast and true streaming formats appear exactly the same to end users, delivery of a streamed podcast is logged the same way as a downloaded file in the server logs. This important distinction impacts the ability to measure content and ad delivery in real-time without access to client side analytics. Podcast publishers must work around this limitation and track metrics using server log data.

Because log-based measurement counts only the file downloads and not the actual listening, in the long run it might not represent the best possible measurement approach. New player-based measurement approaches recently announced by Apple and others show the potential for podcast measurement that's in the client app and can track actual plays, pauses, and other listener behaviors. But it will take some time for any of these client-side approaches to be adopted by all the podcast players that exist, and will depend on the extent of data made available. So for the foreseeable future, log-based measurement is the only way to achieve comprehensive measurement of all podcast usage.

Media delivery via true streaming falls outside of the definition of a "podcast" and is therefore excluded from this document.

Podcast Player Market Share and Tracking Limitations

The ability to track podcast content and ad playback largely depends on the player requesting the file. The native players that operate on iOS systems, namely the Apple Podcasts App and iTunes offer no technology for confirming that a podcast file was played. This lack of client-side response prevents podcast distributors from measuring ad plays at the level expected in other digital media.

In order to provide insight into the limits on tracking podcast content, the IAB Podcast Measurement Working Group was asked to provide reports on the market share for platforms that request podcast files. Podtrac, Blubrry/RawVoice, WideOrbit, Libsyn, and PodcastOne submitted reports used to produce the following table, which aggregates the resulting market share percentages for the month of April 2016.

Aggregate Report on Podcast Player Market Share (April 2016)

Platform requesting podcast file	Range of market share %
iOS - Apple Podcast App	45-52%
iTunes	8-13%
Browsers	6-14%
Stitcher	2-7%
Everything else	12-30%

Reports for non-iOS and iTunes market share varied from one report to another, but the results reported for the iOS Apple Podcast app were consistent across all reports, with a mean of about 49%.

For about half of the podcast ads served to browsers (6-14%), some servers may be able to distinguish ad delivery from a probable ad "play." Using browser plugins or other technology, a specialized tag used to request the ad file can indicate that the player accessed the ad. While this technique offers valuable tracking data, half of the 6-14% means that only about 3-7% of podcast ads can be tracked this way.

Another key insight for this report (not represented in the table) is that less than 3% of the market share enables client-side tracking as it exists in other forms of digital advertising. Only one of the five participants reported a small percentage (3%) of market share for host-branded players (players owned by the podcast producers). Since the producer controls the player and the content, they can request ads and trigger tracking beacons based on ad play. Most podcast distributors see almost zero activity in this market. Even of the 3% of reported host-branded players, many may not be equipped to find tracking beacons and use them.

The data provided for this report shows that half of the market share for podcast players belongs to the Apple Podcasts app, which prevents any client-side tracking or even the ability to count a "play." Podcast distributors must turn to server log analysis and report on ad delivery. Some distributors count an ad once it's been served. This count offers a valuable metric, but is out of scope for this document because an ad served doesn't indicate whether the ad file was downloaded.

Despite the limitations, podcast audiences are growing and offer valuable exposure for marketers. In order to offer this value to buyers, metrics must be consistently defined across the industry. IAB collaborated with members in the podcasting community to establish metric definitions that can be used consistently in the podcast marketplace.

Establishing consensus and clarity for podcast reporting metrics improves communication and establishes trust and accountability with buyers.

Scope

This document defines content, ad, and audience metrics in the context of downloaded podcasts whether saved for later listening or listened to while being downloaded. In this context, both formats are typically pre-recorded and available on demand whenever the listener is ready to access the files.

Podcasts that use true streaming technology to deliver the ad offer the ability to track activity in real-time or near real-time and one metric used to measure "client-confirmed ad delivery" is covered in this document. However, the percentage of market share for applications that support true streaming is currently too small to account for any meaningful campaign measurement. Additional measurement guidance for true streaming audio is covered in the MRC Audio Measurement Guidelines currently in development as of the release of this document. Those guidelines do not cover Podcast measurement – mainly due to the limitations around client side measurement in the industry.

Ad measurement in podcasting presents the industry with many challenges. For the sake of establishing common ground in tracking podcast ads, the definitions presented in this document address counting ad delivery. This count comes from analyzing server log files to determine what was actually delivered.

3. The Podcast Medium – Content Delivery

Podcast listeners acquire podcast files in one of two ways: either by downloading the file for later listening (downloaded), or by listening while the file is downloaded (online listening). To a lesser degree, some podcasts may also be played while a persistent connection to the server is maintained (streamed), but the market share for applications that support this format is insignificant for campaign measurement and excluded from discussion here.

Delivery methods for downloaded files, whether listened to later or during download, offer valuable inventory to advertisers, but content and ad delivery are handled differently in both environments. An overview of each format is explained below. Despite different tracking capabilities in each environment, a few baseline metrics should be able to offer similar reports for both podcast types.

Downloaded Podcasts

Podcast downloading allows the audience to download full episodes of content that can be played at a later date and time. Listeners may subscribe to select programs, and platforms like iTunes continue to support full downloads to a personal library for listening offline at any time in the future. The convenience of this system makes downloaded podcasts a continued preference among listeners.

Online Podcasts

Online podcasts appear to be streamed, but the file is actually being downloaded while the listener is listening to the file. The downloaded file is stored in a temporary location rather than to a library as with a downloaded podcast. Since online files are typically downloaded the same way as the files stored for later listening, delivery for these two formats are recorded the same way in the server logs. The only difference between the two are whether the listener is actively playing the file as it's downloaded or being saved for later listening - which can only be discerned by the player.

Raw Server Logs

In a downloaded file, segments of the file are collected on the listener's device, or progressively downloaded. These progressively downloaded files result in a server log with several requests to the server, which must then be analyzed and filtered from other server requests in order to represent how many files were downloaded and to what audiences. When podcast publishers use a consistent process, metrics can be reported and trusted with a higher level of confidence.

4. The Podcast Medium – Ad Delivery

Podcast ads can be delivered and tracked in a variety of different ways, but in general two different methods are used with variations on each.

Integrated Ads

Historically podcasting ad campaigns often involve ads that are read by the podcast host or a familiar voice. A static ad or jingle may be also included as part of the file. These ads are part of the content and included, or "baked-in," with the file that is downloaded. Targeting is limited because everyone who downloads the file gets the same ads. When the ad is added to the podcast at the time of request, it is dynamically inserted, as described below.

Dynamically Inserted Ads

In recent years, ad technology has allowed for ads to be targeted and dynamically inserted at the time of file request. The ad server determines the best ad to serve to the listener at the time of request. In a podcast consumed online, ads may be inserted into a file that is being progressively downloaded at designated ad breaks. Some publishers may count this dynamic ad serve as an "impression" without confirming ad delivery. The metrics in this document focus on confirming that the ad was delivered. Server logs can confirm that the entire ad file was downloaded, but the process for counting a served ad can only determine that ad file was sent.

5. Podcast Measurement

In digital display advertising, ad tracking is performed using beacons that are triggered in the web browser, or client, which verifies that the ad was viewed or at least had an opportunity to be viewed. In podcasting, client-side tracking is usually only possible when the client player passes tracking data back to podcast producer or distributor. In this set-up, the player is programmed to notify the server when an ad has been played. While this set-up offers the most accurate ad delivery counts, it represents a very small percentage of the podcast industry–less than 3% according to member reports on market share in the industry (see table 1).

Measurement with Server Logs

In order to produce accurate counts for podcast downloads and ads, technical staff in publisher and podcast ad operations must analyze server logs. These server logs may include file requests for a combination of downloaded podcast files, dynamically inserted ads, and any content requested by the web page or application hosting the player. A number of factors are used to analyze log files.

HTTP GET requests, and under certain circumstances HTTP HEAD requests, may be processed that contain the following data.
- **IP Address** - The IP address is one of the factors that may be used to determine if the request is unique or a duplicate. (Exceptions are shared locations such as corporate offices, dorms etc., that have a large number of people sharing the external IP Address) It may also be used to determine geographical information of the media consumer. This applies to both IPv4 and IPv6 IP formats.
- **Time Stamp** - The date and time may be used to determine if the request should be counted.
- **HTTP Status Code** - The appropriate HTTP status code is examined to determine if the request should be counted.
- **Bytes Served** - The value may be used to determine if the media was completely downloaded or if not, how much was downloaded. (Note: This information is only available from native server log files.)
- **Referer** - The origin of the download may be used to determine if the request should be counted. e.g. media that is auto played upon loading a web page may be removed or reported.
- **User Agent** - The identifier of the application or service consuming the media may be analyzed to determine if the request is unique.
- **Byte Range** - The range of bytes requested in a given request may be used to determine what portion of the media is requested.

When analyzed across multiple requests, the information may offer data that represent podcast downloads, audience and ad delivery. Since media technology is always changing, no specific combination of factors or techniques will offer the most accurate count indefinitely. However, meeting some minimum requirements and following some best practices will help produce more

consistent results. The next section will go over some best practices for generating server side metrics.

The document then defines a few metrics for podcast content measurement, audience measurement and ad measurement. Podcast producers and distributors may include additional metrics beyond the ones defined here, but such additional metrics should be labeled separately from the core list of metrics described in this document.

Recommended Process for Measurement

This section lists best practices based on the experiences of the members of the Podcast Technical Working Group. While we have made the effort to be specific, publishers and distributors will have to look at the various options available and select the best for their particular circumstances. To be compliant with these guidelines, the metrics provider should support the process below or a process with a similar or more stringent level of analysis sophistication, disclose the options selected, disclose where they diverge from the recommendations, and provide the rationale/circumstances that drove those decisions.

We recommend a 4-step process to generating metrics using server side log analysis.

1. Apply filtering logic
2. Apply file threshold logic
3. Generate metrics
4. Audit the process (feedback loop)

Step 1. Filtering

All requests that should not be counted for any reason should be filtered out up front. The criteria we have identified for filtering are listed below.

1. **Identifying Uniques (for Downloads & Users):**

    Identifying unique requests is important in counting downloads for an episode and in counting audience size. The following method is recommended, and the details of the filtering methods should be kept transparent.

    Filtering using IP address + User Agent

    - A combination of IP Address and UA is used to identify unique users and downloads. For example, if the same file is downloaded 10 times by 6 user agents behind one IP address, that would count as 6 users and 6 downloads.
    - This method requires a blacklist maintained to block IP addresses that excessively download/play at a value that is not feasible.

- In order to better support known high density IP Addresses (dorms, corporations, etc), a whitelist of IP Addresses may be maintained. For these IP Addresses, it is not necessary to filter by UA.

2. Eliminate Pre-Load Requests

Pre-loading of podcasts directly results in podcast downloads being counted when they should not. There are two possible solutions to handle this.

1. Policy put in place to not allow pre-loading in players (e.g. preload=none for HTML5)
2. Use a download threshold based on ID3 header payload plus 1 minute of recording time to determine if request was for a play/download or for pre-loading (see Step 2 **"Apply file threshold levels"** below)

3. Eliminate Potential Bots and Bogus Requests

There are a number of scenarios where the raw requests include requests that should not be counted because they likely come from bots or from products that behave in ways that make them look like real downloads. We recommend that metrics providers look for the following to filter out.

1. IP addresses that cannot originate to actual users (for e.g., known servers)
2. IP addresses that account for a large number of downloads should be examined for potential fraud. (But also look at the safe IP addresses note below.)
3. IP address that are on a service like AWS.
4. Erroneous referrer data
5. Bogus user agents, e.g. Firefox 3.06
6. User Agents that identify to be from sources that are not actual users (e.g. bots that self-identify as being bots)
7. Similarly, referer data that implies that the sources are not actual users.
8. Apple clients – ignore requests from the podcast app UA, but count the UA for the iOS media player. (This effectively also filters out the 0-1 byte requests from Apple).

Note - Known "safe" IP Addresses (dorms, corporations etc.) should be maintained in a whitelist and be allowed through. These likely need to be re-validated frequently (say every 30/90 days) since IP addresses may not be static.

Note – The members of the Podcast Technical Working Group also suggested that we should start a project to share a whitelist of "safe" IP addresses that represent dorms/corporations and other known NAT situations. Similarly, we likely need a blacklist of known bad IP ranges. We will look into these after publication of this document.

4. Play-Pause-Play Scenarios

If a unique download is divided into multiple file requests, for example if a user plays the first half of an episode using a website audio player, clicks pause, and then resumes a half-hour later, then that should still be counted as one unique download. Care should be taken to not count these as multiple downloads/users.

5. Handling HTTP Requests

This section covers the correct handling of the logs based on the various types of HTTP requests.

1. HEAD requests - these should not be counted unless verifiable that they could be counted via proxies
2. GET requests –
    a. 200 (ok request) should be counted
    b. 206 (partial request) A partial request should only be counted if the range starts at 0. Other partial requests that start at different positions of the media file should not be counted as additional downloads/plays. Furthermore, the final partial request will be logged within most CDN and web servers as a completed 200 request.
    c. 304 (not modified request) -> signal that user has existing file and wants to see if it changed. Only count these if source is determined to be to a proxy not an actual user.
3. There may also be platform specific quirks to watch for. For example, Akamai uses a HTTP code of 000 for 206 requests that ended prematurely.

Step 2. Apply File Threshold Levels

Downloads below a certain size are unlikely to result in human consumption because too little of the file was received to listen to any content. The following rules help eliminate the downloads that are too small to be counted.

1. To count as a valid download, the ID3 tag plus enough of the podcast content to play for 1 minute should have been downloaded.
2. ID3 size recommendation – since the ID3 file size various quite significantly, each publisher should measure the ID3 tag file size for each podcast. To be more efficient in cases where the ID3 size doesn't change, the publisher could set a size for the show/program and whenever the artwork changes, re-calculate the size.

3. Content size recommendation – the size of the download for 1 minute of content will vary based on a number of factors. So our recommendation is for the publisher to calculate this size for each show.

This does require a continuous monitoring of the podcasts as each episode gets served.

Note – 1 minute was chosen as a conservative minimum size since other mediums use similar or smaller thresholds.

Note – if the podcast size is smaller than a minute, the metrics provider can indicate that fact and use the actual full podcast size as the threshold.

Step 3. Generate Metrics

Once the requests have gone through the filtering process above, it is time to generate the metrics defined below, as well as any additional / custom metrics supported. We will not prescribe how the metrics should be formatted / delivered or recommend any particular analytics technology over the other.

Step 4. Audit the Process

The goal of this section is to allow for adjustments to the metrics generation. We need to watch for behavior that indicates that the quality of the metrics is diminished, and investigate the source of potential errors / fraud.

We recommend that the entire process should be self-audited on an ongoing basis. Red flags should be identified and metrics adjusted based on a deeper investigation of the red flags. In addition, future runs of metrics generation should factor in any learnings from each run. For example, if certain IP addresses are identified as being potential bad actors, those should be eliminated in the metrics of the current run and should also be placed in a black list for future rounds.

A good method to audit metrics would be to compare the bandwidth usage to the metrics. The deltas should be linear.

Podcast Content Metric Definitions

Since podcast ads are so closely integrated with podcast content, metrics that measure content are vital to ad measurement in podcasting.

Show producers, executives, marketing, and digital product teams are interested in the following questions:

- Audience: How many people are downloading my network/show/episode?
- Downloads: How many times is my network/show/episode downloaded and potentially listened to, at least in part?

The following metrics are used to describe content downloads. Server log analysis for content delivery should filter data to produce metrics as defined below.

1. **Download:** a unique file request that was downloaded, in accordance with the filtering rules and minimum size threshold rules described earlier.

Once we introduce the "Confirmed" plays from player side listener tracking, we likely can reuse the same set of metrics and add the term "Confirmed" to them, showing both total and confirmed side by side.

Publishers can opt to also provide non-filtered (non-unique) file request information to offer a gross consumption metric.

Podcast Audience Metric Definitions

Podcast listeners usually download more than one episode and often from more than one podcast. A measure of how many people downloaded episodes can be useful by describing the reach of the podcast or group of podcasts.

Website metrics of reach usually depend on the use of cookies or JavaScript to identify individual devices or users. Downloading an MP3 file does not usually trigger the reading of cookies or the running of JavaScript, so alternative methods must be used. HTTP requests - include the IP address of the recipient of the file and usually include a user agent that, while not unique to the user, provides some ability to distinguish multiple users behind one IP address.

2. **Listener:** data that represents a single user who downloads content (for immediate or delayed consumption) from the podcast publisher. Listeners may be represented by IP address (but limited to a max of 3 per IP) or may be represented by a combination of IP address and User Agent as described earlier.

The listeners must be specified within a stated time frame (day, week, month, etc.).

Note – if the metrics provider is able to identify listeners by cookies, IFA or other similar mechanism that requires more advanced client support, they are allowed to do so. The key requirement is to be transparent about the methodology used.

Podcast Ad Metric Definitions

The following metrics represent the first step toward improved ad measurement in podcast advertising. These metrics are derived using the content metrics defined above. As these metrics become adopted in the industry, additional steps can be made toward an improved podcasting ecosystem.

3. **Partial File Download:** a unique file request that was less that 100% downloaded. This metric may be used to help determine ad delivery for ads that were included in the downloaded portion of the ad for files that were not completely downloaded.

4. **Ad Delivered:** an ad that was delivered as determined by server logs that show either all bytes of the ad file were sent or the bytes representing the portion of the podcast file containing the ad file was downloaded.

   For example, if an ad was included within the first 25% of a podcast and at least 25% of the podcast file was downloaded, then the ad can be counted as delivered.

   When ads are dynamically inserted into the podcast file or within an ad break within the podcast, 100% of the ad content (all bytes) must be downloaded before it may be counted as delivered.

5. **Client-Confirmed Ad Play:** counts an ad that was able to prompt a tracking beacon from the client when the file was played. Whenever possible, metric should include information about how much of the ad was played using the markers: ad start, first quartile (25%), midpoint (50%), third quartile (75%), and complete (100%).

   While the client-confirmed ad play metric represents the most accurate count for ad plays in a podcast, it requires client-side tracking. As discussed earlier, the platforms used to download, store, and play podcast files lack or prevent the technology

needed for client-side counting. Aggregate reports on player market share among podcast publishers estimate that less than 3% of players are capable of providing client-side tracking data (see table 1). The IAB Tech Lab will continue working with the player platforms to get more access to client side beacons.

Higher Level Metrics

All the metrics in this document so far were written with a focus on episode level analysis. The Content and Ad metrics described above should also be made available at 3 levels – publisher / show / episode.

Client side tracking or access to some sort of listener ID – for example using cookies or using an Identifier For Advertising (IFA) – would be the ideal mechanism to track audiences over time, but this can be done using analytical methods as described below. If the publishers / distributors have the ability to identify the user, they should indicate the mechanism used and provide the metrics (downloads, listeners and any other additional metrics supported) at the show and publisher level in order to provide a better picture of that podcast's reach.

Lacking a listener ID, there are a few options to build these metrics:

1. Sum the metrics across podcast episodes.
    - This might be ok if the goal is to just count the total number of unique downloads, but does not provide a view of the listener base
2. Use the IP Address & UA to identify and track users
    - This provides a better view of the user base, but it is important to understand the limitations.
        i. The IP Address for a user could change – especially in the case of a mobile user – at which point there is no way to correlate the 2 IP addresses
        ii. Using UA helps differentiate multiple users from an IP address, but breaks down when multiple users from an IP have the same kind of device/UA.
    - However, these two negatives likely counteract each other over time.

6. Summary

Podcast audiences represent a growing segment of marketable media, and are considered to be some of the most engaged. However, the medium is asynchronous and in most cases severed from data collection once delivered, which presents advertisers with measurement challenges. In addition, measurement practices have been fragmented and ad-hoc so far. This document, with the six metrics defined above, offers the first step in an improved environment where buyers and sellers can start to use the same language with clearly defined meaning. As communication improves, producers will be able to scale their operations and invest in technology that brings tracking closer to the standards available in other media options.

7. Appendix: Publisher Player Recommendations

As we developed the podcast metrics guidelines, we recognized that it might be valuable to identify some recommendations that publishers factor in when they build their player / listener experiences.

1. Do not implement Auto-play. This will result in a bad user experience for the user with audio they were not expecting to hear.
2. Do not Pre-load - unless the intent was clearly to play the podcast.

In addition to user experience issues and slowing down the websites, if these guidelines are not followed, measurement companies may decide to discount ALL the traffic from these apps/sites because they cannot count true downloads or plays.