Digital Content Labels & Sensitive Topic Classifiers: Description of Methodology

Last Updated: 2018-11-08

Digital content labels and Sensitive topic classifiers label websites, videos and mobile apps according to the brand safety of the content. You can target inventory that meets or exceeds your brand safety requirements by using brand safety targeting. This technology is available in Campaign Manager (CM), Display & Video 360 (DV360) and Google Ads.

How content is classified into digital content labels

Content is analyzed by Google classification technology and placed in one of five brand safety levels (with DL-NR as the default). This classification process looks at hundreds of aspects of a site's content. Our brand safety classification is updated regularly to reflect any changes and updates to a site.

Descriptions of the digital content labels:

DL-G: Content suitable for general audiences.

DL-Y: Content suitable for families.

DL-PG: Content suitable for most audiences with parental guidance.

DL-T: Content suitable for teen and older audiences.

DL-MA: Content suitable only for mature audiences.

DL-NR: Content that has not been labeled.

Sensitive Topic Classifiers for Brand Safety

For advertisers that need more granular control of brand safety, Google also offers Sensitive Topic Classifiers. These classifiers can be used in addition to or instead of the digital content labels. These are provided for advertisers whose brand safety needs are not uniform across all types of content. (For example, a pharmaceutical company may have stronger concerns about placing ads adjacent to content about drugs and alcohol)

Use of Digital Content Labels in Campaign Manager

Digital Content Labels are used within Campaign Manager (CM) for reporting purposes. From a reporting perspective, users can see a breakdown of the Digital Content Labels associated with the placements that their campaign ran on. For example, a user can see that their campaign ran on 54% placements that were DL-G, 22% labeled as DL-PG, etc.

Reporting on Digital Content Labels is only provided in aggregate - not at an individual URL level.

Use of Digital Content Labels outside of Campaign Manager

The same core Digital Content Label technology is used in a variety of Google products, including:

- Google Ads (pre-bid ad blocking on Google Display Network and YouTube)
- Display & Video 360 (pre-bid ad blocking and reporting)

- Campaign Manager (reporting)

Sensitive Topic Classifiers Implementation

Whereas Digital Content Labels are consistently applied across all Google products, the implementations of Sensitive Topic Classifiers vary slightly based on the needs of each product. The breakdown is as follows:

Google Ads

The following Sensitive Topic Classifiers are available for pre-bid blocking (aka exclusion):

- Tragedy and conflict
- Sensitive social issues
- Sensational and Shocking
- Profanity and rough language
- Sexually suggestive
- Parked domains
- In-game
- Live streaming video
- In-video

Display & Video 360

The following Sensitive Topic Classifiers are available for reporting and pre-bid blocking:

- Adult
- Derogatory
- Downloads & Sharing
- Weapons
- Gambling
- Violence
- Suggestive
- Profanity
- Alcohol
- Drugs
- Tobacco
- Politics
- Religion
- Tragedy
- Transportation Accidents
- Sensitive social issues

Campaign Manager

The same 16 Sensitive Topic Classifiers available in Display & Video 360 are also available in Campaign Manager for both reporting and serve-time blocking.

Further description of each sensitive topic classifier can be found in Google Ads online help guide, located at: https://support.google.com/adwords/answer/3306596

Assigning Digital Content Labels and Sensitive Topic Classifiers to URLs

Each ad impression will have a set of metadata associated with it. This metadata will include, among other things, the URL of the page where the ad appeared[1]. Once the URL is obtained, the content of that URL is gathered via a web crawler. The content of this page is then analyzed for the purposes of assigning digital content labels and sensitive topic classifiers.

A Digital Content Label is assigned to each individual URL. This means that multiple URLs within a given domain can have different digital content labels. For example, one URL on the site xyz.com that deals with adult content may be given a label of DL-MA. However, other URLs on that same xyz.com domain may have a label of DL-PG.

The same URL may also have a sensitive topic classier associated with it. A URL can have multiple sensitive topic classifiers (e.g. the content can be about both drugs and weapons). Conversely, a URL may not have any sensitive topic classifiers associated with it.

Assigning Digital Content Labels and Sensitive Topic Classifiers to Mobile Apps

The process for assigning a digital content label to mobile apps is very similar to that of URLs. The two major differences are as follows:

1. Instead of having a URL associated with the ad impression, a mobile app unique identifier will be present. This unique identifier maps to a specific mobile app.

2. Whereas multiple URLs within a single domain can have different digital content labels, that is not the case with mobile apps. Every ad impression coming from a mobile app will be given the same digital content label and sensitive topic classifiers.
3. Apps associated with a web presence may pass an applicable URL that represents the screen content (e.g., NY Times article). In this case, the web content will be used for brand safety classification.

Digital Content Labels and Sensitive Topic Classifiers Over Time

The Digital Content Label and sensitive topic classifier that is associated with a URL or mobile app may change over time. This can happen for one of two reasons:

1. The content associated with the URL or mobile app changes

2. The technology and classifiers that are used to assign a digital content label are periodically updated to reflect emerging trends, new benchmark data and improved classification technology.

The reclassification process is completed multiple times per day for web content. Mobile apps are reviewed when a new binary is published in the app store.

From both an ad blocking and reporting standpoint, the digital content label associated with a URL is determined based on the latest available content. For example, assume that an ad was served on a URL with a label of DL-T on Monday. If the rating of the URL changed to DL-PG on Tuesday (for whatever reason), the reporting will continue to count this as an impression served on DL-T content.

For exact details on how a date/time stamp is associated with each impression, please refer to the DCM DOM.

Digital Content Label Coverage

Over 80% of ad impressions registered by DCM will have an associated Digital Content Label.

---

[1] *The exact mechanics of how a URL is associated with an ad request is outside the scope of this document.*

The vast majority of cases where a digital content label is not associated with an ad impression is due to the fact that no URL was provided. This could be because of technical reasons (i.e. the ad code on the publisher page is implemented in a way that it makes it impossible to identify the underlying URL) or contractual reasons. For more information on each of these scenarios, please refer to the DCM DOM.

If a URL is associated with an ad impression, hundreds of separate data signals about the domain and the content of the URL itself are analyzed to determine the Digital Content Label. Not all of these signals are necessary to produce a final digital content label. The technology was designed to gracefully handle partial and ambiguous data signals. However, if there is not enough information present to make a final and complete decision, no Digital Content Label will be returned. This last scenario occurs for less than 1% of all URLs.

Below are additional reasons why a URL or app may not have a digital content label or sensitive topic classifier assigned:

- **A site is relatively new**: Sometimes when a site is new, we don't have enough data on the site to classify it. In these cases, the site will normally receive a digital content label one month later.

- **A site's URL is "masked":** Sometimes publishers don't make the URL of their impressions transparent, and in these cases we cannot label these publishers' URLs.

- **Semi-transparent inventory:** "Semi-transparent" inventory may receive any digital content label if the publisher selling the inventory has shared all of the URLs behind the bid request URL with Display & Video 360. In these situations, the publisher's inventory will be labeled according to the inventory's minimum brand safeness.

  For example, if the semi-transparent site semi-transparent-site-123.com has inventory from the following URLs:

  > abc.com (DL-G)
  > cbs.com (DL-G)
  > anothersite.com (DL-T)

  Then semi-transparent-site-123.com will be labeled as being DL-T, since its inventory has a minimum brand safeness of "DL-T".

Please note, inventory is categorized to only one digital content label, and digital content labels are mutually exclusive. For example, this means if you target DL-PG, Display & Video 360 will bid only on DL-PG inventory, and won't bid on DL-G or DL-T inventory. You can and should target multiple digital content labels in the same campaign.

On the other hand, inventory may be categorized to multiple sensitive topic classifiers at the same time.

Digital Content Label Reliability and Quality Assurance
Digital Content Labels and Sensitive Topic Classifiers use a sample of tens of thousands of URLs that have been manually rated by a group of professional content labelers for benchmark data. Every time that the technology which assigns a Digital Content Label is modified, an analysis is performed over this benchmark data to ensure that reliability and performance metrics are met.

For quality assurance on mobile apps, a statistically random sample of labeled apps are further reviewed by a human labeler with more tenure to ensure apps are rated consistently with respect to the published guidelines.

Sometimes ratings are performed by third party vendors who are hired by and conform to the standards of Google but are not Google employees.

If a user would like to dispute a rating they have found in CM, they can click the "Report Classification" button within the Content view ([help center instructions](#)).

Data Retention

Google retains data for each URL, app and domain for which it assigns a digital content label or sensitive topic classifier. At the URL and app level, data is retained for at least 30 days. The following information is retained:

- Unique URL or App ID
- Digital content label assigned
- Sensitive topic classifier assigned

If a URL/app has not served ads after 30 days, the information will be removed from the short-term storage system. For domains, that data is retained for 52 weeks.

In the rare instance when historical information would be needed beyond the dates listed above, the majority of classifier labeling can be found in various log files.

In the case of errors and/or bugs, all issues are logged in Google's internal bug tracking application - known as Buganizer. The data in Buganizer adheres to Google's internal user retention and deletion policy.

Technical Updates and Notifications

Materially relevant changes made to digital content labels and sensitive topic classifiers are noted in our help articles, training materials, and in our client-facing blog posts. Going forward, the DOM will also be updated with product/service change information.

Should any issues arise with our [product/service] which directly affect a client, the customer support team will notify the client of the issue and the possible impact and possible schedule for the resolution of the issue.