

# Implementation Guide for Brand Suitability with the Content Taxonomy 2.2

Version: 1.0

Released for public comment (until November 27 2020)

Please email [taxonomy@iabtechlab.com](mailto:taxonomy@iabtechlab.com) with any feedback

Latest version at <https://iabtechlab.com/standards/content-taxonomy/>

The Taxonomy Working Group consists of members from the following companies:

<ul style="list-style-type: none"><li>• Acxiom</li><li>• AdColony</li><li>• Adform</li><li>• AdGear</li><li>• AdTheorent</li><li>• AlikeAudience</li><li>• Anyclip</li><li>• Audit Bureau of Circulations UK</li><li>• BARC India</li><li>• Browsi</li><li>• Cadent</li><li>• CBS Interactive</li><li>• Centro</li><li>• Signal.io</li><li>• Comscore</li><li>• Confiant</li><li>• Conversant Media</li><li>• Criteo</li><li>• Cyber Communications Inc.</li><li>• Digital Advertising Consortium Inc.</li><li>• Digitas LBI</li><li>• District M</li><li>• DMD Marketing Corp.</li><li>• DoubleVerify</li><li>• Dow Jones &amp; Company (Wall Street Journal)</li><li>• eBay</li><li>• Fiducia DLT Ltd</li><li>• Forensiq</li><li>• FOX Networks Group</li><li>• FreeWheel</li><li>• Fyber</li><li>• Google</li><li>• Grapeshot</li><li>• GroupM</li><li>• GumGum</li><li>• IAB Europe</li><li>• IAB Germany</li><li>• IAB Russia</li><li>• IAB Tech Lab</li><li>• Index Exchange</li><li>• Inmar Intelligence</li><li>• Integral Ad Science</li><li>• Invisibly</li><li>• JW Player</li><li>• KERV Interactive</li></ul>	<ul style="list-style-type: none"><li>• Knitting Media</li><li>• Line</li><li>• Lotame Solutions</li><li>• Macromill, Inc.</li><li>• Magnite</li><li>• Media.net</li><li>• MGID</li><li>• Mobilewalla</li><li>• NBCUniversal</li><li>• News Corporation</li><li>• News UK</li><li>• NumberEight</li><li>• NYIAX</li><li>• OpenX</li><li>• Oracle Data Cloud</li><li>• Outbrain</li><li>• Pandora</li><li>• Permutive</li><li>• Powerinbox</li><li>• Privacy Co-op</li><li>• Protected Media</li><li>• Publicis Groupe</li><li>• Publicis Media</li><li>• Publicis Sapient</li><li>• Publishers Clearing House</li><li>• SoundCloud</li><li>• Sovrn</li><li>• StackAdapt</li><li>• Starcom Worldwide</li><li>• Taboola</li><li>• Terragon Group</li><li>• The Media Trust Company</li><li>• The New York Times Company</li><li>• TikTok</li><li>• TripleLift</li><li>• Triton Digital</li><li>• Trustworthy Accountability Group</li><li>• Unity</li><li>• Verizon Media</li><li>• ViacomCBS</li><li>• VONWERSCH DIGITAL STRATEGIES GMBH</li><li>• Yomedia Network</li><li>• Zenith Media</li></ul>
--	--

The IAB Tech Lead for this project was Amit Shetty, Sr. Director, Product

## Table of Contents:

1. Content Taxonomy overview	3
2. Using the Content Taxonomy	3
3. Brand Safety - Floor & Suitability support in 2.2	5
4. Content Taxonomy usage guidance for buyers	6
5. Implementation guidance for ad verification vendors	7
6. Implementation guidance for DSPs	8
7. Implementation guidance for Publishers & SSPs	9

# 1. Content Taxonomy overview

The Tech Lab [Content Taxonomy 2.x](#) provides a “common language” that all parties (publishers, SSPs, DSPs, verification vendors, advertisers) can use and understand when describing the content of a page, app, or other user environment. The content taxonomy is useful in two main use cases - contextual targeting and brand safety/suitability.

In addition to the Content Taxonomy, the Taxonomy working group has also defined an “Ad Product Taxonomy” and an “Audience Taxonomy” (to describe the product being advertised and an audience segment, respectively). The relevant taxonomies should be used based on the use case.

# 2. Using the Content Taxonomy

The 2.x content taxonomy includes two parts - a set of categories that describes the topic context or “aboutness”, and an additional set of orthogonal content attributes such as content language, format, language, source, media type, etc. These are all associated with IDs that are used when communicating information about a piece of content.

The IDs within the Taxonomy specs should be used when tagging content. The IDs are alphanumeric strings (though a large number are currently sequential numbers to maintain backwards compatibility) associated with each category or orthogonal attribute. The implementations involved may be different depending on application (OpenRTB or VAST or proprietary APIs), but the most common usage is to associate a piece of content with an array of IDs.

With OpenRTB / AdCOM, the “cat” attribute should be used to transmit a list of categories associated with the content and the cattax attribute should be set to 2 (for Content Taxonomy version 2.x). These attributes are available on Ad, Site, App, Publisher, Producer and Content objects.


cat	string array	Array of content categories describing the ad using IDs from the taxonomy indicated in cattax . Implementer should ensure compliance with regional legislation around data usage and sharing.
cattax	integer; default 2	The taxonomy in use for the cat attribute. Refer to <a href="#">List: Category Taxonomies</a> .

For example, in the following hypothetical piece of content -

### 2016 Mini Cooper Convertible First Drive

Sometimes, buying on a budget isn't so bad.

f
🐦
+
✉
👍
in
📌



🖼
📄
⬇
🔍



















Image Credit: Brandon Turkus



**Brandon Turkus**

Conventional gearhead wisdom says to go for the biggest, most powerful engine. For the first two generations of Mini Convertible, this was a no-brainer. You bought the Cooper S. But as Senior Editor Alex Kierstein argued in our [first drive of the Cooper S soft top](#), the less-powerful Cooper Convertible has an ace up its sleeve: a highly entertaining, three-cylinder, turbocharged engine. After some time behind the wheel, this two-time Mini Cooper S (hardtop) owner is ready to say the Cooper Convertible is the droptop Mini you should buy, full stop.

The Cooper's 1.5-liter turbocharged three-cylinder makes just 134 ponies and 162 pound-feet of torque. That's a 55-horsepower deficit and an extra 1.5 seconds, compared to the Cooper S. But who's clocking a Mini Convertible with a stopwatch? The 8.2 seconds it takes to get to 60 mph is perfectly adequate, and the triple's power delivery is additive.

🔧 Engine	Turbo 1.5L, 134
⚡ Power	134 HP / 162 LB-FT
⚙️ Transmission	6-Speed Automatic
⌚ 0-60 Time	8.2 Seconds
🏎️ Top Speed	128 MPH
🚗 Drivetrain	Front-wheel drive
📏 Curb Weight	2,905 LBS
🚙 Seating	2+2
📦 Cargo	7.6 CU-FT (max)
📊 MPG	27 City / 37 Hwy
💰 Suggested Price	\$28,000
💰 As Tested Price	\$37,150

if the content taxonomy matches are  
**Content Categories: Automotive/Convertible (8), Auto Type/Performance Cars(16)**  
**Content Channel: Editorial/Professional (1001)**  
**Content Type: Review (1021)**  
**Content Media Format: Mixed (1026)**

Content Language: en (1068)

Content Source: Professionally Produced (1215)

The OpenRTB/AdCOM snippet on the content object would be:

```
“content”: {  
  ...  
  “cat”: [“8”, “16”, “1001”, “1021”, “1026”, “1068”, “1215”],  
  “cattax”: “2”,  
  ...  
}
```

Similarly, in VAST, the “CONTENTCAT” macro can be used to pass the relevant set of categories as follows:

```
CONTENTCAT = “8”, “16”, “1001”, “1021”, “1026”, “1068”, “1215”
```

*Note 1: The above is just a sample set of categories that can be applied to this piece of content, and intended solely as an example to explain the usage of the taxonomy.*

*Note 2: As in the example above, a piece of content can belong to multiple categories (Convertible, Performance Cars).*

*Note 3: The CONTENTCAT macro is currently being discussed in the Digital Video Technical Working Group, as a replacement for ADCATEGORIES (which was intended to request specific ads, but does not seem useful). We will update this document once the macro gets finalized, in case of changes.*

### 3. Brand Safety - Floor & Suitability support in 2.2

The Brand Safety Floor & Suitability concepts introduced in the 2.2 version of the Content Taxonomy are based on the Brand Safety & Suitability Framework [released in September 2020](#) by GARM in collaboration with the 4As APB. For the purposes of the Content Taxonomy v2.2, the descriptions for each category at each risk level, including the Floor, are described in the framework.

The 11 Brand Safety categories identified in the framework are introduced in v2.2 of the Content Taxonomy as topic categories under the parent category “Sensitive Topics” (id “v9i3On” in the Content Taxonomy 2.2 spreadsheet). The 11 categories are:

1. Adult & Explicit Sexual Content
2. Arms & Ammunition

3. Crime & Harmful acts to individuals and Society and Human Right Violations
4. Death Injury, or Military Conflict
5. Online piracy
6. Hate speech & acts of aggression
7. Obscenity and Profanity
8. Illegal Drugs/Tobacco/eCigarettes/ Vaping/Alcohol
9. Spam or Harmful Content
10. Terrorism
11. Sensitive Social Issues

The risk levels in the Framework that identify the levels of suitability are treated as additional attributes of the content. They are encoded in an orthogonal vector accordingly, allowing “risk” to be associated with a “topic” dynamically. The levels are

1. Floor
2. High Risk
3. Medium Risk
4. Low Risk

Note 1: at the time of release (October 2020), the only topics in the Content Taxonomy that are expected to carry risk associations are the 11 Brand Safety categories. The Tech Lab’s Taxonomy Working Group will work with our members as well as with GARM and IAB to determine whether the risk associations can be applied to other categories in the future.

Note 2: the orthogonal attributes supported by the taxonomy like content type and source could be used as additional signals since they have implications to suitability (for example news).

## 4. Content Taxonomy usage guidance for buyers

### *Brand safety:*

Buyers should familiarize themselves with the GARM/4A’s APB Brand Safety & Suitability Framework and the Tech Lab Content Taxonomy. They should also understand the risk/tolerance levels they are comfortable with by looking at examples of the various suitability examples. They should then work with their DSPs and ad verification vendors to ensure that their goals are met by specifying the risk tolerance goals using the Content Taxonomy.

### *Targeting:*

With the increased awareness of user privacy concerns, contextual targeting is becoming more important. The Content Taxonomy enables buyers to use a consistent language across all publishers and platforms. Buyers should also map any specific areas of interest to the Content Taxonomy categories so that the most relevant content can be targeted for their campaigns.

## 5. Implementation guidance for ad verification vendors

Ad verification vendors are likely to be the main implementers of the brand safety floor and suitability. As such, the recommendations provided here would also apply to any other audience implementing brand safety checks.

Recommendations:

- Ad verification vendors should apply any relevant Content Taxonomy categories to a given piece of content, and/or any of the 11 Brand Safety categories
- Ad Verification vendors should apply a risk level whenever any of the 11 core brand Safety category labels are applied to a piece of content.
- It is possible to have multiple brand safety category labels (and appropriate risk levels with each brand safety category label) associated with a given piece of content.
- As of v2.2, ad verification vendors are not required to apply risk levels to topic categories other than the 11 core Brand Safety categories in the Content Taxonomy.
- As indicated above, the descriptions for the categories at various risk levels are from the GARM/4As APB Framework.
- In addition to brand safety, ad verification vendors might also provide more information about the topic context using the other categories on the content taxonomy.

Since the integration between DSPs and verification vendors is done via proprietary APIs, OpenRTB or other standards are not relevant here, and the guidance provided below should be considered pseudo-code rather than actual samples.

The proprietary API would at a minimum have the following:

1. **Request:** pass in the URL to the content being analyzed.
2. **Response:** return an array of objects, each of which represent a category & suitability for that category. When suitability info is not available, such as when topic categories other than the 11 core brand safety categories are applied, the suitability field would be null.

For example, to describe a piece of content that belongs to 3 content categories (“Arms & Ammunition”, “Hate speech & acts of aggression”, and “Casual games”). Of these, the vendor has determined that the content is at “high risk” for “Arms & Ammunition” and “low risk” for “Hate speech & acts of aggression”, the response could look like the following:

```
{  
...  
"catwithsuitability":
```

```
[
  {
    "category": "avbNf2",
    "suitability": "bsr002"
  },
  {
    "category": "HxqYV1",
    "suitability": "bsr004"
  },
  {
    "category": "693",
    "suitability": null
  },
]
...
}
```

## 6. Implementation guidance for DSPs

### *Brand safety:*

DSPs would likely need to implement the following set of capabilities to support brand safety checks for buyers.

1. A user interface to allow buyers to define their acceptable “risk tolerance” levels (high/med/low) for the 11 core brand safety categories as they set up their campaigns.
2. A workflow where they accept the content URL from the publisher/SSP and pass it along to an ad verification vendor.
3. Use the response from the ad verification vendor (as described above) to decide which pieces of content match the buyer’s brand safety risk tolerance for that brand.

### *Targeting:*

In order to allow buyers to target campaigns based on interest, DSPs should make the categories from the Content Taxonomy available in their campaign creation workflows. DSPs might also have integrations with ad verification vendors to check the content and verify the categories for the content.



## 7. Implementation guidance for Publishers & SSPs

### **Brand safety:**

To support brand safety and suitability, there are 2 key areas Publishers and SSPs should be aware of and implement support for:

#### 1. Floor content:

The key guidance to publishers and SSPs is that they should strongly consider preventing the presence of content that would receive a “Floor” risk level in association with any of the 11 Brand Safety categories.

Barring that, at a minimum they should consider not monetizing or allowing advertising on any such content.

#### 2. Providing context

When making ad requests, publishers and SSPs should tag each piece of content with the relevant content categories from the Content Taxonomy, and provide those in the OpenRTB or VAST bid request (as described in section 2). It is very likely that the buy side platforms will want to perform their own analysis using ad verification vendors of the content categories, as well as the floor & suitability checks. So Publishers/SSPs should also provide the URL of the content, so that the ad verification vendors can perform their checks.

At this time OpenRTB does not support the ability to pass suitability information per category (a general suitability level can be passed now). We are working on an OpenRTB extensions to support that capability. The extension would allow an array of objects with brand safety category and associated risk tolerance levels - similar to the sample in the Section 5.

*Note: We would like to solicit feedback during public comment on whether publishers/SSPs plan to send suitability information in ad requests, and also whether DSPs/buyers would use that signal if it comes from publishers/SSPs. At a minimum this information could be a starting point that can be verified by the buy side.*

### **Targeting:**

Publishers and SSPs should populate ad requests in OpenRTB or VAST (or other integrations) with relevant Content Taxonomy categories for each piece of content, so that buyers can execute contextual targeting on the requests. In addition, the pageurl should be sent along so that ad verification vendors and DSPs can confirm the categorization themselves.