System Metadata Model for Connected Television

Draft A
28 May 2010

Confidentiality and Intellectual Property Notice

The contents of this document are strictly confidential and have been made available to you as a member of the DTG. You are prohibited from distributing, circulating or otherwise sharing this document with any individual, group or company that is not a member of the DTG. The BBC reserves all right, title and interest in this document, its content and subject matter. If you give the BBC suggestions, comments and other feedback you agree that the BBC may freely use, disclose, reproduce, licence and distribute such feedback as it sees fit.

Table of contents

1 Introduction

Descriptive metadata is information describing media content provided by a platform to assist users in the process of discovering and selecting content on that platform. At the highest level of abstraction, descriptive metadata enables three principal content discovery activities by end users:

- **Browsing**.
- **Searching**.
- **Recommendation**.

This specification describes an architectural framework to support these three activities.

1.1 Metadata typology

The two main purposes of descriptive metadata are:

1. To attract end users to media content.
2. To help end users locate the media content and consume or otherwise acquire it.

The richer the metadata available, the more useful it is to end users and the more likely they are to consume content. There is therefore a strong incentive for content providers to publish metadata of a good quality.

In this specification the term **Content metadata** is used to describe the first type of metadata. This is the information describing the media content, such as its title and synopsis, hierarchical groupings, the names of contributors and so on. The Content metadata is, for the most part, invariant in the sense that it doesn't change very much over time as accompanying media content is published. The same Content metadata is often applicable irrespective of whether the media content is being published through a linear channel or into an on-demand catalogue. It makes sense, therefore, to factor out Content metadata so that the same information can be reused in many different publication contexts.

The term **Publication metadata** is applied to the second type of metadata. This is concerned with the temporal aspects of publishing media content, in particular the schedule of events on a linear channel or the availability of a content item in an on-demand catalogue. The linear schedule provides an end user with the information to either tune in to the linear channel at the right time to consume some content, or else book a recording of that content for time-shifted consumption. For the on-demand catalogue, the metadata describes a window of time in which the media content is available to the end user. This too may be consumed immediately or downloaded for later consumption.

This separation of concerns allows the same content metadata to be reused in many different publishing contexts. For example, the same listing can be reused for every repeat of a programme in the linear schedule as well as one or more entries in the on-demand catalogue. The general approach supports a journey in which the end user first discovers content and then decides how to consume it. Equally, it also supports journeys where the user has a particular method of consumption in mind and then selects content available by that means.

1.2 Basic assumptions

The system metadata model makes the following assumptions about the overall ecosystem:

1. Media content is being consumed on a consumer device by an end user. The content discovery and selection process may be occurring through interaction with the same consumer device, but the user interaction could equally be via a different device.
2. The consumer device is intermittently connected to one or more networks through which it is able to receive metadata. At any given time it may be connected to a source of broadcast metadata, an online source of metadata or both.

   The consumer device needs to be able to operate in any of these modes, and it also needs to be able to operate when entirely disconnected from all networks.

3. The consumer device has persistent mass storage that is capable of storing a bounded volume of media content and its accompanying metadata for consumption in disconnected mode.
4. The consumer device will have sufficient resources to **cache** a working set of metadata as a performance optimisation. This may involve pre-emptive retrieval and caching of metadata that may be useful in the near future based on expected user behaviour. It is not practical for the consumer device to persist metadata relating to *all* content on the platform because the volumes involved are too large.
5. The consumer device has a software architecture that supports the execution of multiple different **applications** provided by the platform and by third parties (not necessarily simultaneously). Any of these applications is able to present media content to the end user, including audio, video and still images.
6. One of these applications is a platform-managed **content guide** that provides browsing, searching and recommendation features across a range of content from one or more content providers.

In particular, no assumptions are made about the hosting and distribution of media content. This may be supplied to the consumer device directly from a content provider's own content hosting infrastructure or via an intermediate content distribution partner.

2 Metadata environmental model

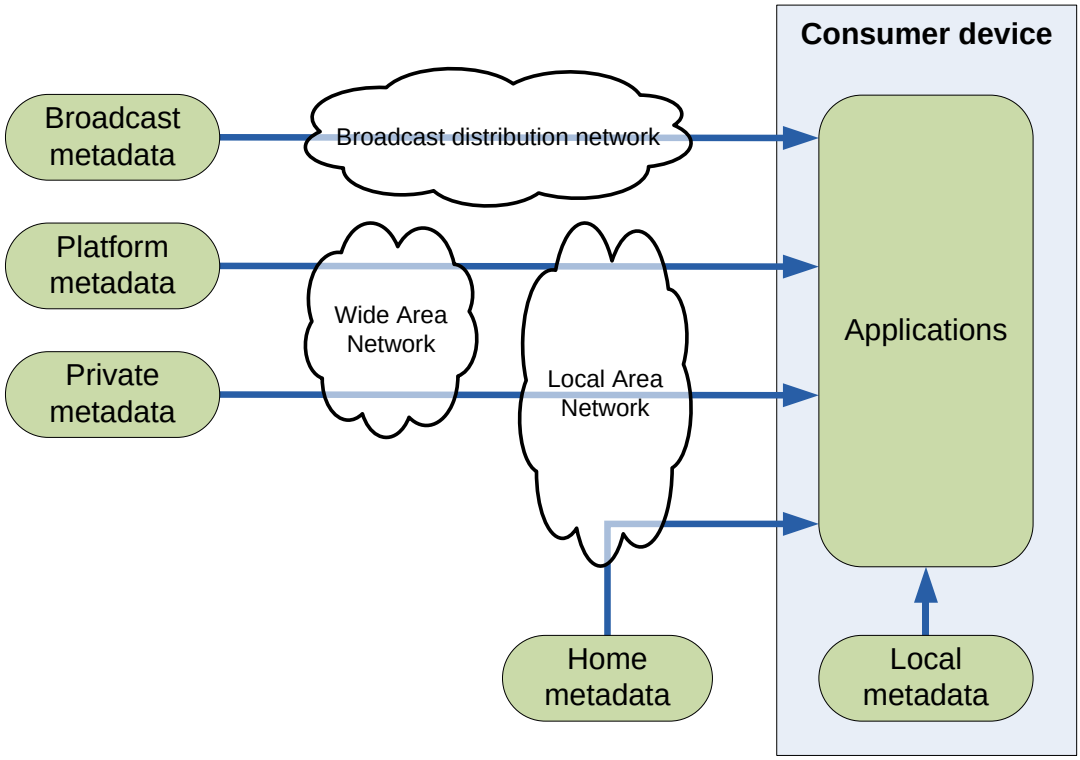

Figure 1: Metadata environment

Figure 1 above illustrates the metadata environment which the consumer device inhabits. Metadata is available from a number of sources, both internal and external.

2.1 Broadcast metadata

This is the metadata received from traditional broadcast networks, such as Freeview or Freesat. It comprises a schedule of linear event listings, and real-time signalling to support the accurate acquisition of content from the broadcast stream. Broadcast metadata is limited in scope and depth by bit rate constraints on broadcast networks.

> *The interface used to deliver broadcast metadata is already well defined in DVB specifications and profiled by the UK D-Book. No substantial changes to this interface are assumed by this specification*.

The principal consumer of broadcast metadata on the consumer device is the platform-managed electronic programme guide application, but third-party applications will also be free to present schedules or other listings information to end users.

The consumer device will be designed in a modular fashion so that different broadcast metadata decoders can be integrated while offering a consistent programmers' interface to applications.

Broadcast metadata is described further in Section 3.

2.2 Platform metadata

This is the metadata targeting platform-managed applications running on the consumer device. The main consumer of platform metadata is the platform-managed content guide application, but other platform-managed applications may also need access to this metadata. Because platform metadata is closely coupled with the user interface it needs flexibility to evolve over time as requirements change.

> *The interface used to deliver platform metadata to platform-managed applications is not exposed to third-party applications and therefore does not need to be standardised.*

Platform metadata is described in more detail in Section 4.

2.3 Private metadata

This metadata feeds third-party applications running on the consumer device. Third-party applications are at liberty to communicate with their own private server-side component(s) using private metadata formats to obtain whatever metadata they require.

> *The interface used to deliver private metadata is out of scope of this specification because it is private to each individual application.*

Again, the need for an optimal coupling between server-side components and the user interface requirements of the application precludes standardisation of private metadata. However, the consumer device needs to provide a sufficiently comprehensive toolkit to support the acquisition, parsing and exploitation of private metadata by third-party applications.

Private metadata is not discussed further in the present specification.

2.4 Local metadata

The consumer device is able to persist metadata for content items recorded or downloaded by the consumer device to its own local mass storage. This metadata supports discovery of local media content even when the consumer device is disconnected from its network(s).

> *This interface is internal to the consumer device and hence private to the consumer device implementation.*

A particular implementation may choose to expose local metadata to applications via an Application Programmers' Interface.

Local metadata is not discussed further in the present specification.

2.5 Home metadata

This is metadata describing items of media content stored on devices within the home. Through appropriate interfaces, the consumer device can discover this content by browsing and/or searching.

Home metadata will be the subject of future work. It is not discussed further in the present specification.

3 Broadcast metadata

The broadcast path is primarily intended to convey schedule metadata for broadcast linear channels. This path typically exists already in any deployment scenario. In the United Kingdom, centralised schedule collation systems are operated on behalf of Digital Multiplex Operators Limited (for the Freeview digital terrestrial network) and Freesat Limited (a direct-to-home digital satellite network). These collation systems provide fully managed electronic programme guide services for their respective platform owners, including management of the linear channel line-up, billed schedule listings and real-time triggers to support accurate content acquisition by digital video recorder receiver devices.

The richness of linear schedule metadata is typically restricted by the limited available bit rate capacity on broadcast networks. The metadata is intended to drive simple electronic programme guide applications; advanced signalling features such as lengthy synopses, cast lists and images are often omitted from deployment profiles to ensure fast data acquisition by simple receiver devices. On the other hand, the metadata collation systems deployed are optimised for low processing latency and real-time performance for content acquisition signalling so linear schedule metadata remains a strong foundation on which to build.

Content in the linear schedule is identified by a non-unique slot number (the **event identifier**). Optionally, broadcasters can also decorate linear schedule slots with an additional **content identifier** that is the same for all transmission instances of the same programme content. This content identifier can be used for cross-referencing platform metadata to establish equivalence between events in the linear schedule and content items in the on-demand catalogue.

4 Platform metadata

A centralised set of platform metadata services is envisaged to support platform-managed applications running on the consumer device. These privileged applications will retrieve metadata from one (or more) platform metadata service providers via the wide area network. Content providers will contribute content and/or publication metadata to one (or more) platform metadata services.

Platform metadata will be both richer and more comprehensive than broadcast metadata. It is primarily intended to support the **browsing** and **searching** of an **on-demand catalogue** and an enhanced **linear programme guide**. The platform metadata services may also be able to process metadata into personalised **recommendations** which can be browsed by the consumer device. These three principal client metadata functions are shown on the right hand side of Figure 2 below.

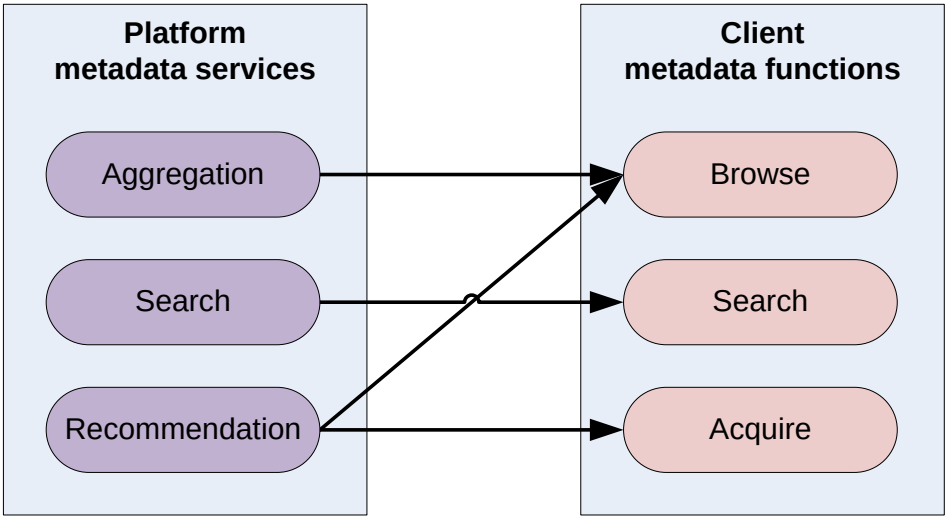

Figure 2: Platform metadata services mapped onto client metadata functions

The three principal platform metadata services offered by the centralised platform metadata services to support the client metadata functions are illustrated on the left-hand side of Figure 2 and elaborated below:

1. **Aggregated Browse**. A set of different faceted views onto the sea of available content items. This includes navigation of the static programme hierarchy (brand and series groupings) as well as a tree of user-facing content categories. Service-oriented cuts of the available content are also envisaged.

   Each such view may be subject to further filtering (e.g. show only free items, or exclude all adult content) and may be available in a number of different sort orderings (e.g. alphabetical, reverse chronological).

2. **Search**. With a large sea of available content, directed browsing is only of use when the end user has an open-ended content discovery goal in mind (e.g. find a political thriller). If the user wants to find something specific, a search-based approach is likely to yield a set of results closer to what is in the user's mind.

   The system metadata model envisages a centralised platform **search service** that indexes all contributed platform metadata. This can then be queried online when the consumer device is operating in connected mode.

   For reasons of privacy, it is unlikely that a centralised search service will be able to index content in the consumer device's local mass storage, or stored in the end user's local area network. A separate client-based search facility is therefore required in addition to the centralised platform search service.

   As an adjunct to the platform search service a **search suggestions service** is also envisaged. This will maintain a list of common search terms and allow a feature whereby a number of likely search terms are suggested to the end user after entering the first few characters. The process of providing search suggestions may be provided as an online platform service, but could also be optimised by caching the list of suggestions in the consumer device for reasons of improved performance.

3. **Recommendation**. Using some combination of explicit editorial recommendations provided by content providers, usage tracking data obtained anonymously from the consumer device and specific recommendations addressed from one end user to another, it is possible to construct an additional set of metadata views of recommended content. These views are exposed by the centralised **recommendations service**. If the underlying platform supports the concept of individual user accounts, these recommendations could be highly personalised (albeit anonymously calculated).

   Irrespective of whether recommendations have universal, household or personal scope, the recommendations service can be used for network bandwidth mitigation purposes. A subset of recommended content items that the end user is most likely to consume is acquired speculatively by the consumer device at reduced cost to the distributor (e.g. by recording the content from a linear channel or by downloading from the wide area network during off-peak usage periods). The media content thus acquired is committed to local mass storage for a limited period and is then available for immediate playback without network impact when the content item selected from anywhere in the content guide, including from the recommendations list.[1]

All three platform metadata services are exposed to the client over a common technical interface that is understood by the platform-managed applications. By sharing a common metadata representation across this interface, the task of handling metadata in the consumer device is greatly simplified. A common approach to accessing subsets of metadata (result set "pagination") can also be shared across the platform-managed applications.

[1] This feature is superficially similar to "broadcast record lists" described in the DTG D-Book, the main difference here being that the end user does not explicitly subscribe to a particular editorial selection of content and the content items acquired are not listed in a user-facing content library. The end user is also unaware that media content is being played back from local mass storage. Also, because the mechanism for delivering the acquisition list is over a two-way network there is more scope for tailoring the list individually to maximise the network bandwidth mitigation benefit. Broadcast record lists and bandwidth mitigation features can co-exist in the same product.

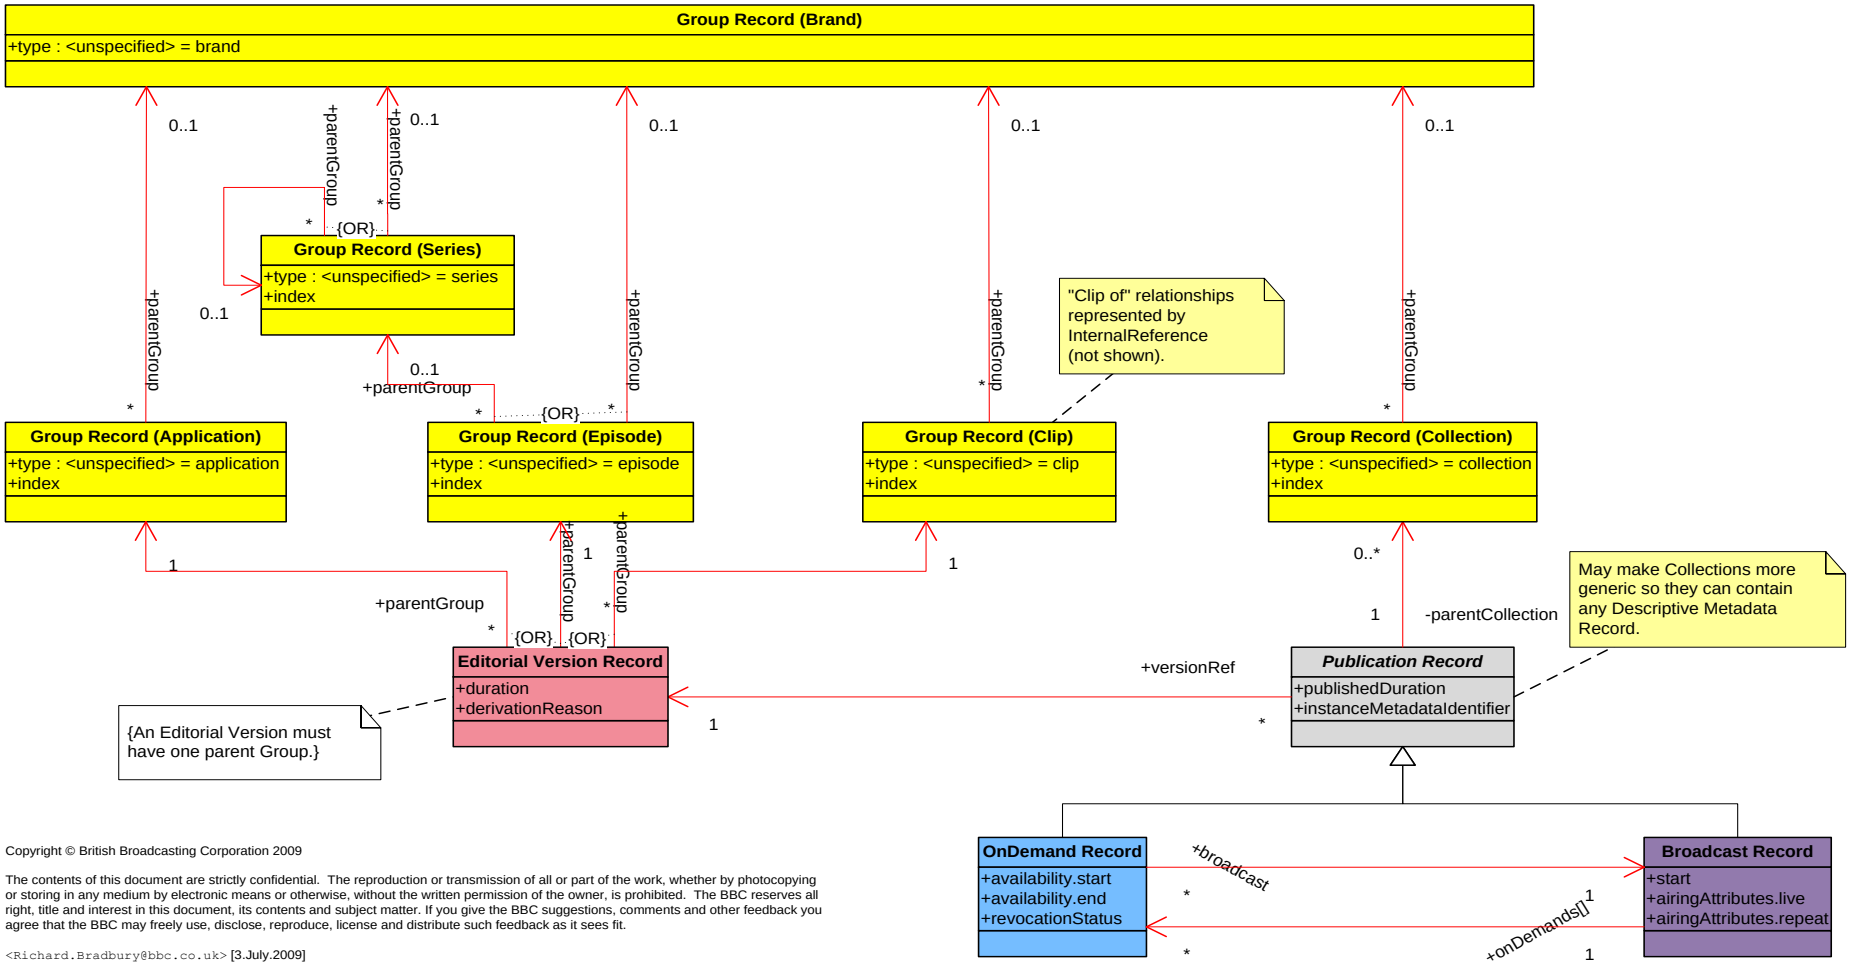

Figure 3: Content metadata in the Logical Metadata Model

5 Logical Metadata Model

A summary of the Logical Metadata Model is shown in Figure 3 above in the form of a UML class diagram. The purpose of the model is to represent all logical entities that exist in the system metadata model. This includes metadata inside the various system blocks and in transit between them. By mapping these entities from the logical domain onto the real metadata systems and interfaces in the ecosystem it is possible to trace individual metadata fields through the end-to-end data flow and identify gaps. The logical model and the mappings are a valuable tool.

The logical metadata model distinguishes between first class entities, which are called **Records**, and other entities, which are referred to simply as Entities. The distinction between the two is that Records are the smallest granularity of atomic update and transfer. Records may aggregate a number of other complex metadata sub-entities, but these cannot exist in transit outside the context of an enclosing Record.

5.1 Descriptive metadata

The majority of Records in the logical model derive from the **Descriptive Metadata Record** (not shown in the figure). All of these Records share a number of common sub-entities, principally Description. Records of this type are further specialised as follows:

- The **Content Record** (not shown in the figure), used to represent the static aspects of descriptive metadata, such as programme grouping hierarchies (represented by the Group Record specialisation) and individual content items (represented by the Editorial Version Record specialisation). Content Records are distinguished by having a content identifier that is unique within the ecosystem.[1]
- The **Publication Record**, used to represent the publication instances of content in a linear schedule (the **Broadcast Record**) or in an on-demand catalogue (the **OnDemand Record**).

  Publication Records always refer to a particular Editorial Version Record by means of its content identifier. They additionally require an instance identifier to distinguish one publication instance of a given piece of content from another.
- The **Service Record** (not shown in the figure), used to represent service-based attribution of content and/or publications. Content Records can point to a particular service to denote a static attribution of the content to a branded service. Individual publications can also reference a service.

  For linear publications a Service Record represents the linear channel: the set of Broadcast Records with the same service reference in a particular time window is a linear schedule for that channel.

  Similarly, the set of OnDemand Records with the same service reference represents an emergent aggregation corresponding to a particular on-demand service offering.

All of these Record types can have broadly the same descriptive metadata provided for them, including titles, synopses, classifiers, credits. They can also cross-reference each other.

There is only one Group Record defined in the logical model, but programme grouping hierarchies tend in practice to be formed from common building blocks: **episodes**, **series** and **brands**. These different types of grouping are catered for by means of a type field on the Group Record. Flexible tree-based hierarchies can be represented because a Group Record may optionally have a parent Group Record. (External business rules need to be applied to prevent nonsensical arrangements of the different Group Record types.)

Two additional types of Group Record are also modelled. The **application** type is used to model the different versions of an application running on the device. Applications are thus on a peer with episodes in the Logical Metadata Model and can be browsed, searched and recommended in the same manner. As with episodes, it is specific Editorial Version Records that appear as item entries in

[1] An example of such an identifier is the globally unique TV-Anytime Content Referencing Identifier (CRID).

the user experience and so an application may have a number of different versions present simultaneously, each with a different window of availability.

Finally, the **collection** type is provided as a means of collecting together individual publications for the purposes of discovery in the user interface. This could be used, for example, to model an upcoming season of events in the linear schedule or a thematic selection of on-demand publication instances.

5.2 Content classification and placement

The tree of content categories exposed to the end user is determined by some editorial function and will inevitably evolve over time to reflect changing editorial policies. The challenge for any system aggregating metadata from multiple content providers is integrating disparate sets of content into a coherent offering.

The simplistic approach to content placement would be to require metadata publishers to label every published content item explicitly with the user-facing category or categories where it is to appear. The problem with this (from the metadata publisher's point of view) is that as the number of target publishing platforms increases, so too does the task of explicitly placing content items into each different category set.

Instead of this, individual content providers typically classify their content according to a private controlled vocabulary.

- In some cases, this is very small and has a one-to-one mapping with the the user-facing categories in the content provider's private on-demand service. Clearly, these are likely to be different from the user-facing categories in any downstream aggregated service.
- At the other extreme, some content providers employ a team of classification professionals to apply a rich set of controlled classification terms to content items. These fine-grained classifiers can then be mapped onto a number of different user-facing categories.

In general, a balance needs to be struck so that content providers with differing resources have fair, reasonable and non-discriminatory access to the aggregated platform.

The approach taken by the system metadata model is to use abstract content classifiers on inbound metadata contribution interfaces and to offer a user-facing category browsing view on the client-facing distribution interface.

Content providers are required to classify their content on a number of different abstract axes such as subject matter (e.g. history, sport, politics) and intended audience (e.g. children, adults). A common controlled vocabulary is targeted by all content providers by mapping their internal classifiers onto the common set. This target controlled vocabulary is hierarchical in nature so that smaller content providers with fewer resources can produce a coarse-grained mapping while larger content providers can produce finer-grained mappings.

The task of the centralised aggregation function is then one of placing content items into user-facing categories according to a mutable set of editorially-controlled business rules. The advantage of this approach is that when the user-facing categories change, only the placement business rules are affected and the impact on upstream content providers is minimised.

As well as providing a point of articulation upstream, the approach can also give downstream flexibility. If, for example, the aggregated platform wishes to expose another set of user-facing categories (e.g. specialised for consumer devices with smaller screens), the same set of content item classifiers can be mapped onto a second set of user-facing categories without requiring any upstream changes.

