SUPPORTING LARGE SCALE VOD THROUGH COMMON METADATA

Yasser Syed, Ph.D.
Cable Television Laboratories

Abstract

This paper explores the challenges in video services and plant operations that the operator has to undergo to support on-demand services. In particular, the paper describe how, a common set of metadata tags will be needed to manage, describe, and present programs to the customer. Without a common set of metadata, there will be limits on the size of content catalogues, types of on-demand services, and search tools for these services.

This paper first describes what is metadata and what are some of the types of metadata (Application, Association, Presentation, etc.) that are used today or considered for use. The next section describes how metadata supports an on-demand service in acquiring content, asset management, navigation, and user interfaces through the existing approaches enabled by CableLabs VOD Metadata specifications or through new approaches. Lastly this article describes how common metadata mechanisms can also be used to increase consumer demand for VoD content & services through guide-like interfaces and other approaches.

INTRODUCTION

The cable industry has traditionally been a broadcast-oriented environment. With digital broadcast services, cable has over a hundred channels transmitted 24 hours a day for viewing. In reality, very few of these channels are watched by a sizeable audience for more than a short period of time. Most of the viewers do not pay for viewing content, but pay instead for the availability of a variety of content that can be viewed. When a cable operator decides on a new channel, it must consider whether the target audience of potential viewers is large enough to support the costs of channel acquisition and resources in the cable plant. Increasingly, this becomes a hard decision because newer channels have smaller target audiences or are repurposing content shown on other channels.

Given this, cable operators have been allocating more bandwidth for video-on-demand (VoD) and Subscriber Video-on-Demand (SVoD) services that provide customers a direct way to view content of their own choice. VoD customers, in turn, need access to simple descriptions of the expanding video library content in order to select videos to view on demand.

This paper explores the changes in plant operations that the operator has to undergo to support on-demand services. In particular, we describe how a common set of metadata tags will be necessary to manage, describe, and present programs to the customer. Without a common set of metadata, there will be limits on the size of content catalogues, types of on-demand services, and search tools for these services.

WHAT IS METADATA?

Metadata is descriptive data associated with a content file or application. A content file in this case could be a Moving Pictures Experts Group (MPEG) video file, a still image file, or an audio file. The metadata may vary from merely identifying the package title, to information for populating an electronic program guide (EPG), to providing a complete index of different scenes in a movie, to supplying business rules detailing how the content package may

be displayed, copied or sold. An asset is an identifiable set of metadata plus its associated content file if it exists.

The information from metadata allows cable operators to distribute, manage, track and present the described content or application to customers. The metadata itself is just text that can be used or ignored. The value in metadata is creating and standardizing common and accepted text fields that can be used to build content.

Categories of Metadata

Metadata can be organized in following categories:

Intrinsic Content Metadata — Metadata that is directly associated with only the content and does not change. This information allows for the content to be useable and routable, but does not necessarily dictate how it can be played or presented. It could be encoding, file size, file type, genre, rating or other information. Identifying what metadata belongs in this category allows for the content to be repurposed for different applications and services without retransmitting the entire collection assets.

Non-Intrinsic Content Metadata — Metadata that is associated with only the content that does not usually change, but is not directly required for use of the content. This is helpful but non-critical information that would assist in playing and understanding the content. Some types of metadata that fall in this category are chapter indexing, actors, language and studio.

Asset Management System (AMS) Metadata — Metadata needed to deliver, distribute, identify and place assets within a headend distribution system. With the proper metadata management wrapper for the asset, it does not become necessary for the headend to fully understand the exact content or metadata it is handling, but just the pertinent information for its asset management routines. An important concept at this level is to be able to uniquely identify the asset such that updates to assets after distribution can be feasible. Types of metadata for AMS are Asset_ID, Provider_ID, Asset_Class, Version, and Product_Offering.

Application Specific Metadata — Metadata associated with applications like VoD/Subscription Video-on-Demand (SVoD) and can span more than one asset for information applicable to a collection or package. Application metadata is required in order to put a collection of assets into a service. Examples of this metadata include association metadata, license metadata and presentation metadata. A key item in this category is the ability to perform operations (e.g., close) on the entire collection of related assets. Specific types of metadata for this are license windows, sequence number, series title (e.g., used to form groups of television episodes), and royalty_ID.

Association Metadata — Metadata on how a group of assets are related to each other. This can be expressed as explicit metadata information or implicit structural information through references or associations. It can show that assets are dependents of other assets or shared among a group of assets. For instance, a VoD Title asset can have as dependent assets: a movie asset, poster asset, license asset and preview asset. The concept of unique identifiers for assets allows for relationships to be easily created.

License Metadata — Metadata of contractual nature for the application. Having license metadata asset for a content provider to create one group of assets for many service

providers. It can then update the group of assets with a unique license asset that dictates the contractual agreement for that particular provider. Some examples of license metadata are license window, contract name, display_as_new, movie preview_period, royalty_percent and billing_ID.

Presentation Metadata — Metadata that involves presentation of the video offering to the consumer through a user interface. This information can change often according to display constraints and marketing strategies. Some examples include summary, title_brief, and category. For example, category is a marketing field for the assets that can change over time (e.g., $1 movies, Christmas Movies, Weekend Blockbusters). This is different from genre, which are permanent fields that can categorize the content and can be used for navigation purposes. Other aspects of presentation are display characteristics. Common Metadata formats needs to also consider character limitations, Multilanguage aspects, scripts and symbols that can simplify displaying the information on the user interfaces for all types of STBs.

User Metadata — Metadata that can be used to target content to the right type of consumer. In a passive way, metadata fields in the application and content assets can be used to allow the user to easily identify content they want to see. It can also allow the service provider to target ads for users watching a particular type of content (e.g., advertising a blender for those watching a cooking show). In a more active way, metadata can be created for a user that can indicate preference for particular content. This then can be used by the application or service to display content that the user prefers, and would be more likely to order (e.g., a personal barker channel for the user).

HOW METADATA SUPPORTS AN ON-DEMAND SERVICE

Metadata can be used to support all key processes of an on-demand service and can enhance performance as these services become more heavily used. This section describes how broadcast services use these functions in the headend, and how this usage changes as on-demand services become more popular. Lastly, it shows how metadata could improve the performance of these functions for on-demand services and, in general, headend operations.

Acquiring Content

In broadcast, content is usually a retransmission of a program channel from either a local broadcast (ABC, NBC, CBS) or a premium-branded channel (HBO, Showtime, TNN, MTV). Other types of content that are directly acquired by the headend are local advertisements and pay-per-view movie content, which are placed in ad insertion and video servers. The arrival of this type of content is negotiated in advance and delivered in batches (often by tape and now more frequently by closed satellite link) to the headend in a timely manner.

In a VoD/SVoD system, content, which is typically movies, is similarly negotiated beforehand and delivered in batches to the headend in a timely manner. The catalog of movies is often a magnitude larger than pay-per-view selections, but still limited (~200-300 selections). Content is acquired via tape delivery or over closed satellite link.

In true on-demand services, the consumer expects the equivalent of a personal program channel. To acquire content in a scaleable fashion, content delivery will necessitate the following behaviors in content acquisition:

- It will incorporate different types of content (Movies, TV episodes, Daily news, Sports, Instructional Video);

- It will happen more often because the size of the content catalogue will greatly expand and be actively modified (e.g., Sony's content library contains over 6,500 movies and 35,000 TV episodes![1]);

- It will be distributed in a manner where the complete package will be delivered over time starting with previews, continuing with asset updates, and finally the primary content file;

- It will allow for reuse and repurposing of content for different types of applications and services after initial offering.

The collection of assets can be delivered over a variety and combination of mechanisms including satellite, tape, IP network, and e-mail (see Figure 1). In most types of network delivery systems for assets, there is a "pitcher" and "catcher" function that are aggregation points for transmitting and receiving the assets. The "pitcher" aggregates metadata and content from studios and encoding houses, makes corrections to the assets when required, and usually acts as the business entity responsible for content negotiations and scheduled delivery. The "catcher" is the aggregated receiving point that validates the delivered asset, communicates with the pitcher, and instantiates and distributes the assets internally to multiple headends in the cable network. An important function of the catcher is to validate the delivered assets such that mistakes are not propagated throughout the cable system, and then to communicate with the pitcher in case a "repitch" of the asset is needed.

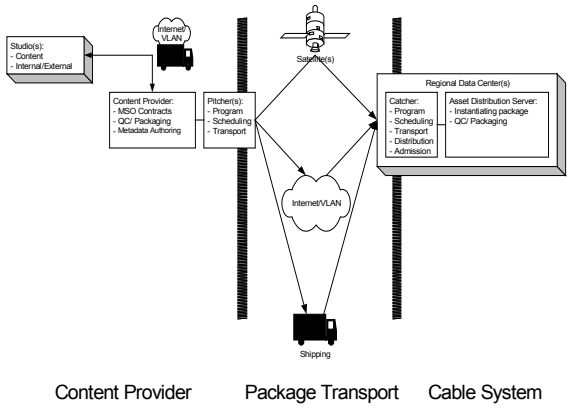

Figure 1: Cable Content Acquisition Process

Metadata can assist in this process by uniquely identifying content assets and re-associating them with intended applications and product offerings. Additionally, it will allow for reassembling the collection of distributed assets such that content (the most time-consuming element in the transfer process) can be sent separately from metadata and other assets.[2] Furthermore, validation and corrections (another very time-consuming process) can be automated.[3] Licensing details for a particular piece of requested content can also be automatic once the overall contract negotiations for services between the content provider and cable operator are established. As on-demand

[1] "Broadband and VoD Come of Age: Sony Exerts its' Enormous Influence", Multichannel News, March 4th, 2002

[2] Metadata is a small fraction of the size of the content file (KB vs. GB). Resending content with metadata whenever a metadata update is required can be a time consuming process especially overlossy networks. Separation of content and metadata can allow for easy updates of metadata while enabling content to be sent over other mechanisms that are more suitable and secure for transferring large files.

[3] It can be very time consuming tracking and determining mistakes in the metadata fields once the assets have been ingested into the system because corrections are usually manually inputted. Savings in time can happen if the asset can be validated before ingestion and rejected. This allows for the content provider to correct the mistakes and repitch.

services grow, the content catalogue also grows and becomes more active on a scale that is magnitudes above what is currently done. Automating and distributing some of these processes to acquire content is absolutely necessary at this scale. Having the appropriate metadata fields allows for automation of these processes.

Asset Management

In broadcast, asset management was limited to simple support of initial set-ups for retransmission and remultiplexing of broadcast channels, pay-per-view servers, and local ad-insertion servers. Broadcast asset management was also limited to determining the program channel line-ups on the plant's physical spectrum. In VoD systems, asset management functions become more complicated because content is stored on distributed servers in the cable plant, and physical spectrum is a resource that is actively managed to service movie requests and peak-usage estimates. Furthermore, the billing system in VoD moves more towards a real-time function that can handle single VoD requests and subscription VoD models.

As on-demand services scale to larger volumes, there will be more application/content servers and distributed storage devices. An elaborate internal distribution system to continually redistribute content and assets via a combination of central distribution and edge-caching servers will need to be supported. The billing and reporting interfaces will need to validate and calculate bills in real-time, taking into account time of day, local marketing offers, subsidized packages (e.g., on-demand with commercials), historical data, order reports, and customer profiles. Lastly, scheduling, determining available plant bandwidth, and call admission strategies will need to be determined on-the-fly in accordance with the volume of requests.

Standardized metadata is absolutely necessary to handle the multitude of asset management and business functions within a headend system on a large scale (see Figure 2). Metadata, such as asset class, can be used to route assets to the appropriate application or content server. Resource allocation algorithms can be developed using metadata fields like "year" to determine whether content should reside at a central server or be allocated to local edge-caching systems. Licensing information can be used to determine availability of content and place this content in service at the appropriate time. Royalty information and user metadata combined with marketing offers would be needed to determine in real-time actual validation and purchase price to the customer. Lastly, for the call admission, or session set-up, bandwidth allocation algorithms that know the runtime of metadata would be helpful in predicting near-term bandwidth capacities. There are numerous other examples. These are just some of the ways a common set of metadata can assist in asset management in the cable plant, especially on a scale that requires automation in its processes.

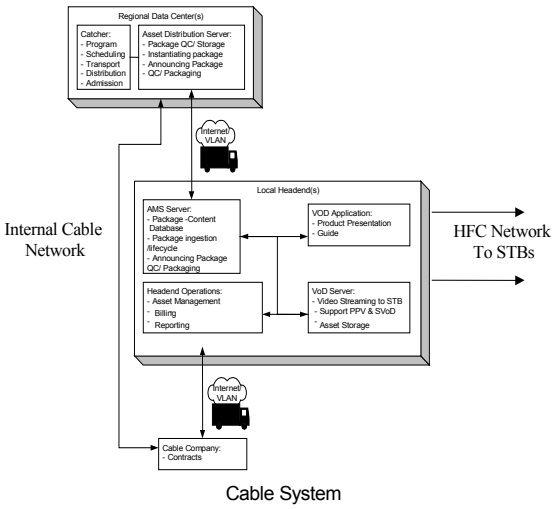

Figure 2: AMS within a Cable Network

Navigation

There is little search capability in the broadcast environment. The EPG serves as the navigator to the consumer. The customer basically looks on the EPG, sees the possible content with some small descriptive information, and the time/channel it will be shown. It is up to the customer to view in person, or to schedule a recording device, when the content is played.

The initial VoD systems are slightly better by allowing an interactive program guide interface. The customer can search for a movie alphabetically and by predetermined categories (e.g., just available, action, comedy, HBO movies). The customer then can see a description of the movie, and maybe the preview, to determine if he likes it. This type of approach can work for limited-size catalogues (under 200 selections) and becomes less friendly as the size of the content catalogue grows.

As the size of the content catalogue and services grow, a better navigation system is necessary. An interactive query device to a database that correlates metadata fields with content becomes necessary. Having the proper metadata fields and elements are required to create a user-friendly search tool for the customer.

With common metadata tags in place, different types of query needs can be easily accommodated with the only restrictions being the type of fields defined, and the way the search tool is set-up. For instance, if someone wanted to watch Spanish-speaking shows to learn more Spanish, they could do a query that would look for country of origin or language option field. In another example, if a customer wanted to watch only Bob Villa instructional videos on building a porch, they could do a combined query on actor, description, application title, and genre. This ability to query a database for a particular request becomes a useful feature as the size of the on-demand content catalogue becomes unmanageable for traditional navigation interfaces.

Presentation Interface

Printed media is an early form of an on-demand service. Books need to be presented in a manner to the customer such that he/she wants to buy it. If someone does not realize a book exists, chances are that book will not be read. But there are numerous presentation mechanisms that make the user aware and want to read that book. This comes in the form of book lists, print and media advertising, and in word–of-mouth.

Broadcast uses similar mechanisms to advertise its own content. Broadcast channels are a good place to promote content since the chances are a sizeable audience may watch a promo for another show while seeing their favorite program. The drawback is there are only a minimal number of spots with sizeable exposure and some of those spots are needed to sell as advertising spots. As the amount of content increases, there is less opportunity to present each one. But in broadcast as a program grows in popularity, it in turn can become a vehicle to advertise for other content because there is a known viewer audience. This is hard to replicate in an on-demand model and could be one of the reasons broadcast will still exist even though an on-demand service may become the preferred way to deliver most content in a cable network.

For present-day VoD, this presentation mechanism is not fully developed. Typically exposure happens through the EPG, a barker channel showing movie previews and, rarely, a spot on the broadcast. The VoD or SVoD services (not content), on the other hand, may

have lots of exposure via local ad insertion on some of the cable broadcast channels. The only constant mechanism that exposes customers to an unknown particular piece of content is the barker channel. As the content catalogue increases, it would be increasingly hard to give the right exposure to content without targeting first the content to the user.

Metadata can be used to present the content to the intended targeted audience. Using the concept of barker channels, category barker channels could be created based upon genre or a designed market offering. For instance, there could be a barker channel for action movies [genre], $1 offerings [category], or what's the latest on HBO SVoD [display_as_new and product_offering]. Looking at more integrated solutions, metadata can be used to insert relevant previews for other content [based on genres] within a content that is currently being viewed by the customer (this is like trailers being shown in movie theatres).

Combining this with user metadata, there are even more dynamic ways of targeting content to the viewer. For instance, user statistics can be utilized to indicate those people who like this movie also liked these following movies (this approach is similar to what Amazon.com or some of the Internet audio-on-demand services do). Another approach is to use the customer's history file, or inputs, to search through a content catalogue and present selected previews that might suit the customer's preference. This can then be presented on the customer's navigation or Graphical User Interface (GUI) screen as he is looking for something to watch.

The key idea is that a common, agreed upon set of metadata can enable targeting of content to the user in a passive or active mannar by presenting the right content for selection to the user. This becomes more valuable as the amount of content prevents the user from making a complete search on his own. In this case, instead of searching, he is presented with a selection of content that might be preferred. This is even more efficient than creating another niche channel in a broadcast environment.

CREATING CONSUMER DEMAND

The mechanisms for creating consumer demand for content in Broadcast are well established. It involves creating program channel brand identity (e.g., American Movie Classics, HBO, Showtime). There are also advertising promo spots in other shows that already have an audience share, as well as advertising in print and media. Lastly, one popular show can lead the audience into another show by simply scheduling it right after (e.g., "Must See" Thursdays on NBC). These various types of methods to create consumer demand capture an audience by:

- Creating an interest in the content through advertising spots that consumers are likely to see;
- Creating a brand for the program channel that will draw an audience that believes any content on this channel will be interesting;
- Using the program order to lead an audience into another piece of content.

Most of the devices to create consumer demand for content in these cases are heavily reliant on the broadcast model.

For an on-demand service, new ways to create consumer demand needs to be developed since a schedule-driven programming approach does not exist. These

new ways will be heavily reliant on making use of metadata for both the content and user. Some possible new concepts are:

Video Magazine — This is like having a video version of a National Geographic, Teen Beat, Rolling Stone, MustSeeNBC. Each magazine would contain pieces of content that would be of interest to users drawn towards the brand. The brand of the magazine would create the consumer demand. Some of the brands or indexing in the brands would be assisted by metadata information.

Personal GUI Interface — The user would have a personal interface that would suggest content of possible interest for viewing. This content can be acquired by user metadata or by user inputs on content metadata like genre. Similar concepts are already being used in services like TiVO. There are several levels of complexity dealing with managing multiple users (family, individuals in family, etc.) as well as developing a proper user model (i.e., avoiding misinterpretations like my TiVO thinks I'm bilingual when I'm really not). Alternatively, the GUI can also be customized based on region/ethnic information of content and user (e.g., show option for a Yankees baseball game in NYC and a Giants Game in San Francisco).

Spot-Insertion — For the right price, even viewers of on-demand content could tolerate promos for other pieces of content while viewing something they ordered. These promos can be selected based upon the metadata of the current content they are viewing. For instance, if a user was watching an action movie, they can be shown promos for other action movies. Promos can also be presented based upon the preferences of the user.

Marketing Offers — This could be similar to a lead-in and could be used for recurring series types of content. For instance, a viewer could order the most recent episode of an up-and-coming series and maybe get to view an unaired episode of a popular series Like the Supranos that the same studio is producing. Creating these offers and personalizing them would necessitate knowledge and use of the metadata fields in lots of creative ways.

Availability Windows — The on-demand service can support previously made content, but could also support content just becoming available (e.g., 1st release of movies, current episodes of TV series, content direct to on-demand). To automate services based upon availability, this will require knowledge of the license metadata as well as metadata related to types of assets.

Broadcast Target-Insertion — The broadcast channel will still hold the largest audience for a single viewing session. The promo spots can be used to promote on-demand content or services. In one approach, the entire audience can view the same promo spot. In a second approach, a different promo spot (but paid by the same studio) can be sent to different viewers based on user metadata and content metadata. A product advertisement can also use a similar approach in both broadcast and on-demand environments.

Search Engines — The ability to search for content in as many intuitive ways as possible is highly desirable. To do this, developing user-friendly search engines based on common metadata fields are necessary. This allows the user to tailor individual requests that can query a database (e.g., a Robert Redford fan can locate all movies he starred in as well as directed). Without a search tool, the only other recourse is to develop a guide based upon general, anticipated types of

requests (this will still also need metadata to keep the guide current).

Subscription On-Demand Services— A subscription service can create demand by allowing users to sample different pieces of content without paying for each individual item. The user feels that he is paying for a service rather than an individual selection. This concept uses similar viewing habits as developed in broadcast where the user 'surfs' the channels for content, and can integrate the users towards trying on-demand services. This can also be a way to promote non-subscription content that needs to be purchased because the subscription service has only a limited choice (e.g., a Sci-Fi subscription service can advertise for the season premier of "Quantum Leap" while offering last year's episodes on its menu).

CONCLUSION

In each of these cases, applications to create consumer demand can be developed around metadata. This process can come from a cable operator, content provider, or third party, but an agreed upon set of common metadata fields are required to develop consistent demand creation applications. An underlying substructure for maintaining constant customer demand is a sizeable content catalogue and service request availability. Common metadata can help in both these back office functions by creating powerful asset management systems that can maintain a large content collection, as well as management of bandwidth resources to minimize denial of service on VoD requests. Lastly, common metadata from both the user and licensing assets would be needed to automate billing functions and data to support these new types of demand-creation opportunities.

REFERENCES

[1] Cablelabs Asset Distribution Interface Specification Version 1.1, www.cablelabs.com/projects/metadata

[2] Cablelabs VoD Content Specification Version 1.1, www.cablelabs.com/projects/metadata

[3] Cablelabs VoD Content Encoding Profiles Specification, www.cablelabs.com/projects/metadata

[4] "Wrestling with On-Demand Services: How to Make VoD a Common Household Name," D. Defrees, CED, Nov.2002.

[5] "Managing On-Demand: The Impending Challenges of Ramped Up Content," Y. Gorden-Kanouff, Communications Technology, Sept. 2002.

[6] A Heavenly Collision: when the HD and VoD worlds come together, will operators be ready?," G Hardin and J. Matarese, CED, Oct. 2003.