

NEXT GENERATION VOD ASSET MANAGEMENT MANAGING THE VOD CHAIN END TO END

Chris Stasi - Vice President, Operations
TVN Entertainment

On Demand product has become a significant source of content to consumers but has only just begun to tap into its potential. As more and more content is accessed, stored and viewed in an on-demand environment, the requirements of asset management become exponentially greater. Systems will need to coordinate asset distribution from the time assets are exposed from a provider throughout their lifecycle, including distribution, ad insertion, streaming and financial reporting. One centralized system, managing content scheduling, offer creation, updates, distribution, tracking, metadata management, edge playout and wholesale reporting makes both MSO and provider VOD systems more efficient and scalable.

This paper will present an asset management strategy for both MSO's and providers covering the entire lifecycle of all pieces of the VOD architecture.

- *MSO asset tracking, updating and reporting*
- *MSO headend management*
- *Distribution management*
- *Content Provider Asset Scheduling, tracking, updating and remote monitoring*

On Demand content has only begun to scratch the surface of what is achievable in the consumer entertainment market. As the nascent technology becomes more and more refined a new level of interoperability and asset management is required. This includes not only the content itself but the associated

files and data as well. Currently the total hours of content available in the On Demand market is on the order of 3,000-4,000 per month. As the platforms become ubiquitous, and the technology easier for the average consumer to understand, the number of hours will increase quickly to 10,000 hours and beyond. Ominously, this does not scale linearly. With the introduction of the new elements and features of the Cablelabs standards and the rapidly growing incursion of advertisers into the On Demand market, the *type* of asset management required changes, as well as the amount.

Where We Are Now

Current asset management consists, predominantly, of two areas: metadata and content files. Both require a somewhat straightforward management style.

In the current Cablelabs¹ environment the metadata file is capable of a relatively small number of functions. It is the steward of the content from the time the content is created/encoded until EPG population. In the interim it guides the asset into residence on the VOD server and populates the associated databases. Subsequently its only functions are updates to certain allowable fields (on capable systems) and deletes.

The content file obviously functions solely to be viewed and the corresponding management of it is currently straightforward. It exists only as part of its own, single, offer/package and when the license window runs out it is deleted from

the video server. It is also tracked solely by references to it within its single metadata file. All of this is about to change and those changes will introduce an enormous layer of complexity into the On Demand environment.

Changes on the Horizon

The biggest issues facing On Demand players are Cablelabs upcoming 2.0² standard and the changing nature of the On Demand content itself. Both cause some of the same issues and many of those issues can be solved in similar ways.

The first asset management issue is the Cablelabs specification change. Whereas currently there is one offer associated with each content file, and each file exists in only its own, single, content package, in 2.0 the content file(s) each exist at their own level. There are one or more title/offer level metadata files tracking each asset and potentially multiple potential files tying each asset to other assets. Thus the single content asset that currently exists on its own with its own metadata file will now be able to sit on a server as part of many different offers and be acted upon in many different ways. This affects the content provider offering the content, the distribution company sending and tracking the file, the server allowing access to the file, the EPG's finding and displaying each offer of the content and the billing systems tracing usage. Many of the links in this particular value chain are not currently capable of the required changes.

The second issue is the changing type of On Demand content. The On Demand platform began mostly as a sister entity – programmatically – to the Pay Per View product. Movies were the first entry as well as Premium channel's SVOD offerings. Consumers were most familiar with this

product and had shown a willing appetite. Most of this programming was long-form and therefore small in number and easily trackable. Over the past 18 months the introduction of music videos, barkers, magazine shows and various other short-form assets has shown the direction content is headed. Advertising is the obvious next step and with its plethora of short-form, unique, content appears to be the 900-pound gorilla entering the room.

Macro Problems

So what asset management problems present themselves in this environment? There are a myriad of small ones, the devil is always in the details, but the major ones can be broken down into some sizable buckets:

1. The lack of a widely used, ubiquitous data set surrounding all On Demand content
2. Vendors within the On Demand space exhibiting different capabilities and requirements
3. Protecting the content to the satisfaction of content providers
4. The inability to update content in a near/real-time, meaningful way

Problem 1: Differing Data Sets

Potentially the biggest threat to widespread content usage is the inability for providers and MSO's to track content in a meaningful way. Currently the Cablelabs specification does a comprehensive job giving all parties the same ability to do so but not everyone is taking advantage of the information. The most widespread current Cablelabs spec is 1.1 and it has a fairly straightforward, well thought-out structure. Each content file, be it the main asset, a corresponding piece of artwork, or a trailer,

has its own asset ID. The metadata file corresponding to that asset and all the offer information it contains also has an asset ID, called Title ID. The combination of all this is the content Package and it has an ID of its own. The media files should keep their ID forever and any changes to the metadata should result in a change to the Title ID. Any change whatsoever to any aspect of the package results in a change to the Package ID. However this is not occurring across the industry. Widespread interpretive differences of the spec as well as embedded functional roadblocks have caused many providers as well as equipment and distribution companies to reach a different conclusion about how the data is managed. This has caused issues along the entire length of the On Demand chain. A provider has problems because if he/she does not track a content file with the same Mpeg ID there is no way to know how many times that same piece of content has been reintroduced or to get any meaningful comparative usage data back from an MSO. A distribution company cannot maintain any semblance of an organized library and distribution system if an ID structure is not followed for the life of an asset. A server vendor cannot reasonably be expected to cover all various types of content introduction, and make any type of meaningful error reporting on content, if it cannot expect a standard set of rules to be followed. If Titanic is introduced into VOD for the first time at one price, which filters throughout the server and billing interfaces, and the following year returns at a lower price, but the Title ID isn't changed, there is no standard way to report on the different price points. Data fields within the metadata files, such as Billing ID's, cause the same issues.

Solution:

So what can be done to make it easier to distribute and track content in an On Demand environment? The easiest and most obvious is for providers to all work off a common data system. Not that the onus should fall only on them but as they are the beginning of the road they get the first responsibility. The Cablelabs 1.1 format allows all providers as much flexibility as needed to identify their content. The use of domain name as the first half of a concatenated ID guarantees uniqueness across providers. The assignment of the ID's *within* each provider is then up to the provider itself to ensure and maintain uniqueness.

Following that, it is up to the distribution companies to maintain the provider's ID integrity throughout the industry. Maintaining this trackable ID structure is crucial to insuring that the content itself is traceable from a providers own management system, updatable from that same system and won't clash with assets being distributed to a location by another distributor.

At that location it is then up to the differing equipment vendors, mostly gateway, server and billing systems, to be able to ingest the ID string in the way it is intended. Currently there is a severe disconnect throughout the industry at this point and it causes not only a loss of traceability but literally a stoppage of content propagation and usage-data flow, which results in revenue loss. Following the protocols in the chart below shows the intended use of each ID level and what should and should not be passed.

	Original Offer	New Offer
Provider	TVN	TVN
Provider ID	TVN.com	TVN.com
Title	Spider Man 2	Spider Man 2
<u>License Start Date</u>	1/1/05	3/1/05
License End Date	1/31/05	3/31/05
Package Asset ID	TVNX1234500000000000	New
Title Asset ID	TVNX1234500000000001	New
Billing ID	56785	New
Movie Asset ID	TVNX1234500000000002	Same
Movie Content Value	TVNX1234500000000002 movie.mpg	Same
Preview Asset ID	TVNX1234500000000003	Same
Preview Content Value	TVNX1234500000000003 tr.mpg	Same
Poster Art Asset ID	TVNX1234500000000004	Same
Poster Art Content Value	TVNX1234500000000004 photo.bmp	Same

Upon the introduction of Cablelabs 2.0, the data set, if possible, becomes even more important. In 2.0 the elements exist on their own, apart from a singular, ever-present package ID. Collections come into play and they enable content assets to be acted upon by metadata that hasn't originated with the asset's initial introduction. A collection could consist solely of metadata introduced halfway through a movie's lifecycle offering a discount if the movie is viewed in tandem with a related new release feature new to the video server. If the collection metadata refers to an asset ID and that ID is not recognized the same way by every system nationwide the collection is not usable and again revenue is lost because of a disconnect on how best to interpret data.

The features and benefits of moving everyone into the same ubiquitous data set are easy to see. The most basic is simply the ability to track where your content is and who is using it. Building on that enables the introduction of ads. If you know where a hit movie resides and everyone knows the same ID structure it's easy to introduce an ad into a collection or playlist available to the end-user. It also enables the main content to be re-used without requiring re-pitching and

reintroduction to the server each time, saving money, time and processing power.

The widespread adaptation of the same data structure everywhere is a hurdle that has to be overcome quickly in order for On Demand content to evolve into the next generation product everyone desires.

Problem 2: Different Requirements from Edge Equipment

Similar to the previous problem, but potentially trickier to solve, is the differing level of capabilities and requirements of hardware throughout the On Demand architecture. These differences can be seen in areas as diverse as encoding specs, data requirements or EPG display capabilities, yet they all have a solution within the realm of asset management.

On Demand Asset management doesn't mean simply tracking data but also the construction and movement of the media itself. This area has encountered some difficulties because of different requirements across different platforms.

A perfect example is that the encoding specifications within Cablelabs³ have set the

baseline for asset construction but there are far more areas for missteps within the grey area of encoder setup itself. A configuration such as Program Stream has a choice within the specification but if an asset encoded on a valid, compliant stream arrives on some servers it will fail to ingest correctly even though it is a “legal” encode. Obviously a spec in a new environment cannot possibly be expected to have uncovered every nuance across all manufacturers’ equipment, but the introduction of elements such as Gig-E switches on the VOD network and other increasingly common architecture changes require a bit more standardization.

As mentioned in the previous section, billing systems currently present a few problems in the next generation On Demand environment. Many are set up to deal with ID’s from the PPV world. This was fine when there were assets numbering in the dozens arriving at a system each month and the only difference was time of day they were being watched. On Demand content has already reached the level of thousands of pieces of content per month and is growing rapidly. No provider is going to allow his or her content to be used forever without meaningful revenue in some form. That requires either direct usage-tracking for paid transactions or click-data for advertisements. Many current billing systems cannot provide this data based off the 1.1 and 2.0 ID structures. If this continues, the introduction of playlists, collections and increased shortform data will come to a standstill.

EPG’s also need to be able to offer the same minimum usage experience regardless of the platform they rest on. Currently each EPG is capable of different data capture and displays. While there will always be commercial reasons to offer different capabilities, a certain baseline set is

beginning to emerge through the Cablelabs group.

Solution:

So what can be done to standardize the requirements for different vendor equipment in the On Demand environment? Tough question. Anyone who has spent any time in the space knows that not only are the normal political elements involved in equipment purchase, vendor relationships and sales incentives but the element of speed to market also comes into play much more heavily in new technology centers. Manufacturers are forced to design on the fly and implement upgrades and rollouts before they’re ready in order not to lose market share to a competitor. But this shouldn’t stop the industry as a whole, through Cablelabs and other like organizations, from beginning to implement standards that go beyond the data and baseline encoding specs that are currently deployed.

While there will always be differences between rival commercial products, which obviously should be encouraged, there is no reason not to ensure certain quantifiable, measurable benchmarks be achieved in more areas. In order to be Cablelabs compliant a server and a distributor need to meet certain criteria in metadata creation and mpeg settings. Other settings need to be implemented as well, such as minimum capabilities on an EPG or minimum data compliance from the billing side. There is no reason a billing system cannot be certified as compliant in the On Demand environment the same way a server vendor must be. As more and more small MSO’s get into the On Demand space, and as more and more new companies bring their equipment into it, the natural propensity will be to become more disparate in the technology sets, not less.

Broadening the standards to include more equipment, encoding, display and interface requirements is the necessary next step to ensure advancement and interoperability.

Problem 3: Content Protection

As a content distributor, it has become obvious over the past few months that one of the biggest waves cresting on the horizon in 2005 is content protection and encryption. Providers are increasingly focused on ensuring their content is as safe as possible as it moves through its digital lifecycle. In an On Demand environment, where content is being passed from studios to distributors, sent over satellites across the country and taken down in hundreds of headends with almost as many differing security protocols, it's easy to see why a major motion picture studio or television network would be concerned about their content falling into the wrong hands. The seemingly endless downward spiral of profits within the music industry due to peer to peer file-sharing and illegal downloading is nothing if not alarming to content owners, and with good reason.

The question obviously isn't *whether* to protect and encrypt content but *how* to do it in a way that makes managing the protection protocols and the content itself as easy and as transferable as possible. Without some format for doing this, managing the asset throughout the VOD chain will not be possible.

The biggest problem implementing such a format is the number of places the content is touched in an On Demand environment in order to reach the end-user. It starts with the dub house, moves to one or more distribution/transmission companies where it is encoded and multicast and ends up at the MSO/Telco's headend. It is there that it

ingests into the video server and resides during its license window for streaming to the customer's set-top box. Every one of these touch points is susceptible to either the easy duplication/theft of content or at least the appearance thereof. The studios will not differentiate between the two and will require all touch points to have the same level of security.

The fact the studios haven't been more focused on this issue so far is due to a myriad of reasons, one of which is the seemingly safe environment of the secure headend. However this won't last. According to a study by AT&T and the University of Pennsylvania,⁴ 77% of trackable movie theft has been traced back to studio insiders. If the actual film companies themselves have been so compromised, an MSO headend, where the files are already digitized, is an even easier mark.

Solution:

So how do we prevent such a potential platform-destroying problem from taking root? The issue, obviously, is to protect the content as thoroughly and as long as possible. But with all the previously mentioned touch points for any single piece of content how does that take shape?

The best and most scalable solution is to encrypt any piece of content from its moment of creation (encoding) and offer the choice to leave that encryption on, or decrypt, at the natural handoff points throughout the distribution chain. With a modular, strippable solution in place the content is never clear-text except in a closed network, while it is simply changing cipher, not eliminating it. So how does this look?

There are multiple types of encryption and no two MSO's will want to implement it exactly the same way. Some will favor pre-encryption and some will favor session-based. Many factors will go into this, each as specific to the MSO, and as valid, as any other. The best way to enable all of them is to offer the content to them in a way that makes either possible.

A distribution company is generally the first player in On Demand to touch the content in a digital format. (This is a generalization as many content providers do their own encoding.) The first thing any distribution company must do is create an MPAA approved library storage system that is approved from a physical access standpoint and also stores encoded files solely on a private network. This obvious step will keep access to the files to as few people as possible.

At this point the files should be encrypted, but how to do that this early and still fit into the modular scheme mentioned above? A standard 128 bit AES encryption scheme can be applied at this, or any, point and passed down the chain. As the content is received at an MSO's headend the receiving device can be configured to either decrypt the content prior to handing to the video server or pass it through in its encrypted state. If the decryption takes place the content must be handed to the video server across a private network to ensure its safety. If the MSO chooses they can receive the content still in an encrypted state and pass it all the way to the set-top. This of course requires the set-tops to host a decrypt client that can receive the decryption keys. It is in the MSO's best interest to get set-top manufacturers to certify as many of these clients as possible to ensure competition and an open standards-based solution.

The encrypted data packets are transparent, so firewalls, proxies etc. can see the encrypted packets and pass them through as if they were unencrypted. This allows any encrypted file to pass itself down through the value chain and be acted upon (transferred, streamed, trick-file creation) as if it were clear text. If this end-to-end pre-encryption is carried through it is up to the distribution company to manage the key-list server and enable all of its MSO's to receive these keys. It also has to ensure that the encryption scheme is standards based. Getting all the encryption companies to maintain a standard scheme, and potentially a universal cipher, is underway and will go a long way to making the transition to a completely encrypted product easier.

If an MSO chooses to receive the content as clear text the obvious choices are session-based encryption, with its higher degree of safety but commensurately higher costs, or to re-encrypt on the server. This requires the same client decryption on the set-top.

There are many complications as encryption enters the On Demand environment. By adhering to a standards-based, modular, flexible encryption scheme that allows the cipher to be removed or passed through at any point, the studios and networks will be satisfied their content will be protected as thoroughly as they demand. This will remove a major roadblock from managing On Demand assets throughout their lifecycle

Problem 4: Immediate Data Access

A problem that has not derailed the rollout of On Demand content but will soon pose a big problem is the inability for programmers, distribution companies and even MSO's to act upon the data surrounding their content in a real-time,

meaningful fashion. As any new technology rolls out to customers, the wow factor carries it through its initial hiccups. This fades quickly and soon the product needs to meet consumers' up to the minute demands or it will fade into obscurity. Once the technology is settled the marketing takes over. The On Demand product chain has been very limited in what the marketers can do because the data they can get their hands on has been very delayed. Even the small amount that comes to them quickly can't be updated or changed in an easy, ubiquitous fashion. This will have to change quickly for the On Demand product to gain relevance to the end-user. It will also be a requirement for advertisers. Nobody wants an ad running that is out of date by the time he/she has the ability to update it.

With the upcoming release of the Cablelabs 2.0 standard the ability to wrap metadata around content differently will help bring some marketing capabilities into the On Demand environment. An advertiser can use either a collection or a playlist to attach ads to free On Demand content or aim targeted ads at a willing viewer. An asset price or window can be changed, or that asset can be packaged with like material to use one as a loss leader, after the material has been in its viewing window for some time. There are still some major hurdles that need to be overcome before this is anything more than a technical capability though. There is no performance data available to providers to determine how, for example, a New Release movie is performing from market to market until the early heavy-usage period has passed. Even with this data at their fingertips the providers have no way of acting upon it remotely without involving another vendor or distribution partner. The only way to aim these ads correctly, or know what offer to adjust on a new title, is to have immediate access to usage data and a system

to take advantage of it.

Solution:

The easiest was to prevent this data restriction from adversely affecting the programming in On Demand is to open up the usage data and give everyone along the distribution chain access to changing it. This consists of two major steps.

The first is to allow a standardized interface into the server usage data. This doesn't necessarily mean giving providers or distributors private customer information. What it does mean is at least passing back generalized buy rates and click data. This can be done through a standardized reverse interface from the server outward. A distributor can link into the server and extract (or more likely be fed) a set of usage data that can filter back to the distributor for storage in their database. This data can then be mined to determine what pieces of content are performing better or worse than expected. Then, assuming the MSO's equipment is updatable (which by CL 2.0 it would have to be) the data can be used to extend, change or refine the offering to the consumer, making the asset more valuable. This interface has begun to get some traction within the Cablelabs consortium.

This information will be worthless however without the ability to view it, make decisions and act upon it. A distributor receiving this information from a server needs the ability to grant any provider using his/her system access to the data and update functionality in real time. A distributor's system needs to be able to author content within it and send that content to the destination headend's but that's not all. In order to make data management meaningful, any data brought back to the provider

through the reverse interface needs to be editable by the same provider without having to go headend by headend. One centralized system, able to update data, nationally or singularly, is necessary. This system needs to manage content preparation, transmission and headend interfacing at every level, in other words an overall layer sitting on top of the entire value chain. This enables any interested party, anywhere on the value chain, access (with permissions of course) to their content, be they provider, distributor or MSO. Without this universal overlay the job of manipulating content will grow exponentially larger and quickly get out of hand. This will have the effect of making content quickly grow stale, thereby seeming less interesting to the consumer. As advertising and short-form assets continue to be a larger part of On Demand content and as the number of files associated with each piece of content grow, one system capable of managing this content from end to end is a necessity.

Moving Forward

On Demand (in one form or another) will eventually be how most people watch most content. Many more issues than discussed here will eventually pop up, some of which take everyone by surprise. But planning for the next generation of asset management starts now. Done right, everyone involved can benefit, from consumers through the smallest distributor.

REFERENCES:

1. <http://www.cablelabs.org/projects/meta-data/specifications/specifications11.html>
2. <http://www.cablelabs.org/projects/meta-data/specifications/specifications20.html>
3. http://www.cablelabs.org/projects/meta-data/specifications/content_encoding.html
4. http://www.aeanet.org/GovernmentAffairs/gamb972_ATTReport_MoviePiracy.asp?bhcp=1