INGEST & METADATA PARTITIONING: *REQUIREMENTS FOR TELEVISION ON DEMAND™*

Robert G. Scheffler, Chief Architect
Broadbus Technologies, Inc.

ABSTRACT

On demand video services, such as today's Video on Demand (VOD), Subscription Video on Demand (SVOD), and the fast-approaching Television on Demand™ (TOD®) are enhancing the consumer television experience and creating new, exciting revenue opportunities and increased cash flow for cable operators and content owners alike. However, the technical requirements to support these services are becoming more demanding and complex. In VOD, cable operators are seeing solid buy-in rates, repeat purchase patterns, and concurrency rates of 3%-10% with limited marketing and promotional support. With recent trials of SVOD and an increased number of popular titles, concurrency rates have 'smoothed' the peak usage rates throughout the week to numbers that often approach 10%-20%. However, with Television on Demand (TOD) services, consumers will have considerably more programming choices including movies, subscription-based content, and the most popular broadcast content. It is anticipated that concurrency rates of TOD may steadily climb to levels that approach 30%-65% -- rates that mirror the total concurrent U.S. television viewing audience as measured by rating services such as Nielson.

Increased service usage, additional content, and new business models are challenging MSOs to conduct unprecedented network architecture preparation and planning. In addition, decisions related to asset distribution, content propagation, network loading, metadata and rules issues need to be addressed to make Television on Demand a commercial reality.

This paper will address the issues and requirements associated with server ingest of broadcast content and content propagation. It will also discuss the architectural implications for the VOD server and propose a new class of server to support TOD requirements. The paper will also discuss how TOD content is managed through the creation and distribution of enhanced metadata formats in an environment that is controlled by studios, distributors, and cable operators.

New video server architectures and rules-based content control and propagation systems become integral contributors to the success of future on-demand services.

VOD/TOD CONTENT INGEST

The issue of the ingest of broadcast television content is one that will become more and more important for advanced video services such as Television on Demand to become a reality. As more content is made available and concurrency rates increase, architectural decisions will have to be made to support these increased demands on the network. A new architecture comprised of higher density VOD/TOD servers with the capability to ingest broadcast television will be required to support ever increasing content libraries and stream counts. However it is important to look at the evolution of VOD architectures to understand how those requirements will change in the future.

VOD in the Past

In the early days of VOD, movies were distributed on tapes. These tapes were shipped to each site that required a specific movie title. Using an encoding rate of 3.375 Mbps and an average movie length of 100 minutes, the total size of each movie was roughly 2.4 GBytes. A typical installation might contain a library of under 100 movies and was capable of streaming to less than 1,000 subscribers simultaneously.

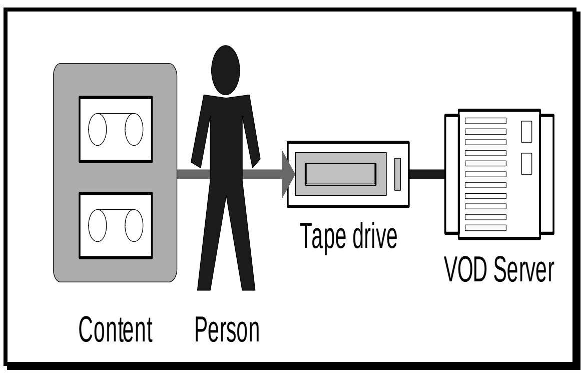

Figure 2-1 Content Ingest for VOD in the past

In early VOD deployments, metadata or other business rules weren't typically supplied with the content. The operators themselves were responsible for deciding what rules applied to particular content and for entering the appropriate rules into the VOD server or control system. This relatively simple model meant that most of the attention was focused on the billing interfaces, set top box (STB) client, and head-end control. With low stream counts, movie titles could be loaded during off-peak hours when the VOD server had more processing capacity to focus on the ingest functions. This was very labor intensive with a single operator feeding tapes and entering rules to instruct the STB guide software about the pricing and availability of new titles (see Figure 2-1). Keeping up with content ingest was quite manageable for the operator and the conventional VOD server.

VOD Today

As an industry, VOD has matured beyond the simplistic example described above. VOD installations now enable 1,000 to 3,000 customers to access a library of 150 to 300 movies. As a result, shipping tapes to VOD enabled head-ends has proved to be a logistical challenge and has evolved to a newer model called pitch-and-catch, where content is distributed by private broadcast to remote stations and syndication partners via satellite (see Figure 2-2). With increased library sizes, increased stream counts and more diverse suppliers sending data, the distribution and propagation of content has shown itself to be quite a challenge. Content can still arrive on tapes and is caught by catchers along with trailers, posters, and rules that are required to put it all together.

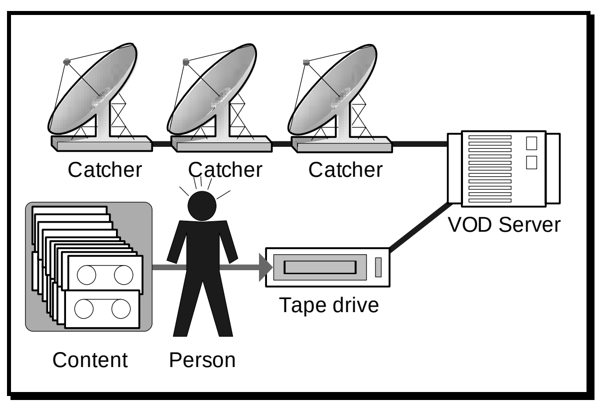

Figure 2-2 Content Ingest for VOD today

Content aggregation companies have risen to the challenge by offering services to edit, adjust, and compile these diverse formats and metadata into a nice bundled package to be pitched and caught. However, a fundamental problem is that while quite adept at low-volume streaming, conventional VOD servers usually lack in their ability to simultaneously ingest large quantities of content. The situation multiplies itself as we add streams, services, storage, and begin to distribute more hardware throughout the network.

Combining SVOD with VOD

Subscription VOD (SVOD) increases the existing VOD content library by adding 50-100 movies and other content and making them available to an increased number of subscribers. Even with a limited amount of content offered, trials of SVOD to date have resulted in increased concurrency rates that may be as high as 10%-20% or 3,000 to 5,000 streams in a typical system.

These concurrency rates place tremendous demands on the streaming capacity of the network. Also as stream counts increase, so does the problem of content ingest. To increase the stream count, additional streaming servers are required. These additional servers need access to the library of ingested content. If a given piece of content is to be made available to every customer on the network, the content needs to be either locally stored or remotely accessible. One way to make the content accessible is to add an ingest server or propagation server at the point where the content is caught or loaded from tape. This ingest server could then locally store the content, making it available to the rest of the servers. Alternatively the ingest server could be used to propagate or distribute the content to the streaming servers, whether local or remote (see Figure 2-3). Remembering that the streaming servers are primarily intended for streaming, there is a fixed amount of bandwidth available for large amounts of content propagation. To now handle the ingest of a significant amount of content, a conventional VOD server will typically lose some, or much of its streaming performance.

Today's VOD server systems adequately accommodate the demands of low-concurrency VOD/SVOD deployments. However, adding the task of ingesting numerous channels of broadcast content to conventional VOD servers creates a massive hardware and software infrastructure that takes up a lot of space, consumes a lot of power, and is inherently less reliable.

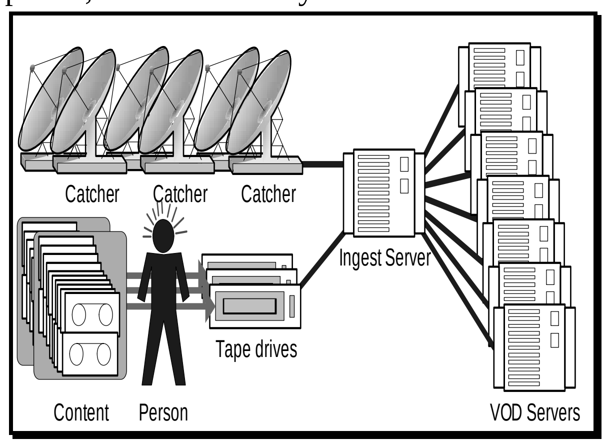

Figure 2-3 Content Ingest for VOD in the future

Television on Demand using Conventional VOD Servers

Now let's look at an example where we expand the VOD/SVOD service offerings to include Television on Demand (TOD). TOD enables cable operators to provide on demand delivery of live or pre-recorded broadcast television services as well as the movie and subscription-based content that VOD/SVOD offers. TOD is especially attractive to television content owners because it allows the viewing and sale of older programming that is out of syndication. TOD enables the consumer to have PVR functionality during broadcast television viewing without requiring a hard-drive in the STB. At a minimum a TOD system should be capable of storing 1,000 movies for VOD/SVOD customers, plus 10,000 hours of captured broadcast television.

With Television on Demand, ingestion, propagation, and streaming of content needs to occur such that the customer still feels like they are watching broadcast television. In addition to the plethora of content, trailers, posters, and rules that VOD/SVOD requires, there is now a real-time requirement for low latency content ingestion. Current VOD/SVOD systems, complete with catchers, tape drives, and content ingest propagation now have to support the ingestion of broadcast television feeds (see Figure 2-4). The path of the broadcast feed to the broadcast ingest server to the ingest and propagation server to the VOD server and then to the customers is an operation that will take many seconds and must occur at the same time as the propagation of VOD content to the VOD servers.

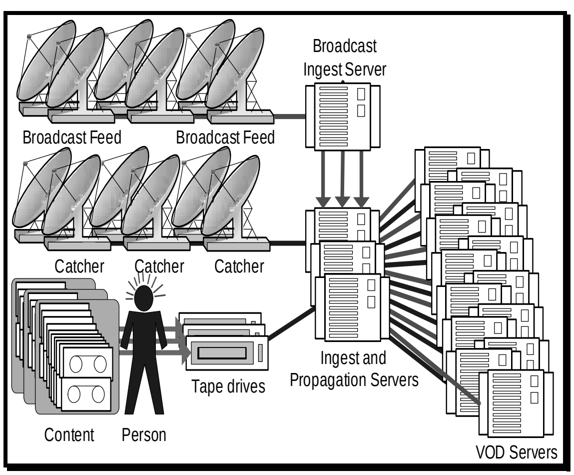

Figure 2-4 Content Ingest for TOD using VOD servers

Consumer concurrency rates for TOD will require a much higher stream count than the current growth projections for VOD offerings. When VOD, SVOD, and TOD are combined a typical system may require 20,000 to 40,000 simultaneous streams. For example, using conventional VOD servers capable of 500 streams each would require 80 servers to satisfy the stream requirement. However, as more conventional VOD servers are added, the problem of propagating the content to all the servers increases exponentially and creates the need for more ingest servers to propagate the content so that eventually there is a hierarchy of ingest servers to streaming servers. A conventional VOD server is designed for streaming to customers, not for moving, propagating and ingesting television content. Therefore today's VOD severs are not the optimum solution for this compelling, new application.

Deploying TOD with TOD servers

The critical issues that must be addressed to adequately support TOD are content ingest and stream count. A new class of TOD server is required that can ingest dozens of channels of broadcast television while simultaneously redistributing thousands of streams with zero-latency. The associated delays can be removed by running the broadcast feeds for ingest directly into a TOD server where they can be directly streamed to customers without requiring an external hierarchy of propagation servers. This solves the content ingest and propagation problem presented by TOD. However, a hierarchical approach to storage is also required for off-line VOD/SVOD/TOD content access. What is needed is a distributed storage strategy with shared local storage as well as shared remote storage that decouples the streaming functions from the storage functions. By decoupling these functions, stream-count and storage-size can be scaled independently while storage can be placed in the network where it can be used in the most cost-effective way. A master head-end containing a pooled storage library would allow a group of servers to access lesser-used programs without requiring local copies. By using this distributed storage architecture, each type of content can actually be moved and positioned in the network for the perfect balance between hardware and transport costs. As the needs of the network change, the placement of system components can change as well.

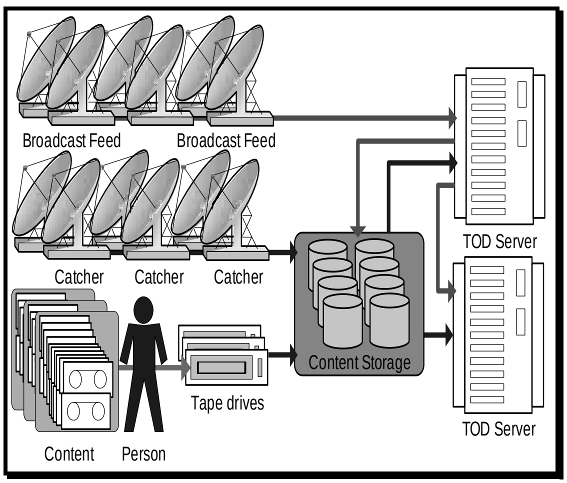

Figure 2-5 Content Ingest for TOD using TOD servers

A flexible architecture that can handle low-latency live-ingest as well as pitcher-catcher and tape based distribution models would be ideal for cost-effectively supporting TOD applications. The capability to decouple streaming from storage, while being able to distribute the storage anywhere in the network, would also significantly improve the economics of TOD. With streaming positioned in one place and storage distributed throughout the network the new architecture will scale to support even the most demanding TOD applications (see Figure 2-5). The future of VOD, SVOD, and TOD are dependent on a new architecture where scale can be controlled and each environment can be tailored for specific applications with unique requirements.

Summary of Content, Streams and Ingest

As needs grow and new business models are introduced, the capacity and scaling of VOD streaming servers are being tested. Content libraries are increasing and greater concurrency is leading to higher and higher stream counts (see Figure 2-6). With the introduction of TOD, the added ability to ingest broadcast television with low-latency is transitioning from an interesting feature into an absolute requirement.

Conventional VOD servers are being taxed to the limit with only a modest library change rate per month. As content libraries grow, to prevent libraries from becoming ‘stale’ with old content, an increased demand is being placed on off-line ingest. Even now, conventional VOD servers are reaching their limits in being able to keep up with SVOD and VOD applications. Regardless of how much streaming requirements increase as TOD begins to proliferate, the cable operator will be forced to add additional servers just to handle ingest tasks. Even then, the resulting system will not adequately address the problem of broadcast ingest to streaming latency. The clearly superior solution is to use a new class of specialized TOD server capable of ingesting and directly streaming with no perceivable delay.

Application	Movie Library	Library Change	Real-Time TV Ingest	Concurrency Rate	Stream Count
VOD	150-300	15/month	0 streams	5%-10%	1,000-3,000
SVOD/VOD	200-400	40/month	0 streams	10%-20%	3,000-6,000
TOD/SVOD VOD	1,000	100/month	100 streams	30%-65%	20,000-40,000

Figure 2-6 System Capacities for VOD, SVOD, and TOD

METADATA AND CONTROL

Rules are needed

The business of broadcast television today is very complex. The participants are numerous -- content owners, content aggregators, content distributors, broadcast and cable networks, MSOs -- and the relationships between the players are dynamic. What keeps content flowing from creators to consumers is the execution and enforcement of detailed contracts. These contracts determine the rules of "how", "when", and "by whom" content may be viewed. Whether it's a re-run episode of "Friends" that airs in syndication on TBS or a live broadcast of the New York Knicks on ESPN 2, there are specific contract-based rules that govern the manner in which content is handled. Therefore, it should be no surprise that a system of contract-based rules will continue to govern (and perhaps with greater emphasis) in a business that combines broadcast television content with on-demand content.

When VOD was initially deployed, the rules were relatively simple. MSOs would license a window of time when a movie would be made available to its subscribers. During the licensing window, the movie would be placed on the VOD Server and be available to subscribers. After the window was over, the movie would be deleted from the server. A set of rules, or metadata, capturing the pre-negotiated License Window Start and End Times would be read and enforced by the VOD server.

As the industry moves towards SVOD and ultimately TOD, the same set of complex rules and attributes must be applied to each piece of content. Examples of additional rules for handling television content could include:

- Specific days of week when content is available
- One or more timeslots during the day
- Time range that the program is available on a particular day
- Specific commercials that must be carried with the program
- Trick-mode rules and attributes (specific speeds, enabled/disabled functions)
- Specific customer groups by demographic or geographic regions

Rules should be entered and applied as early in the process as possible. There are rules from many levels. Examples include:

- Content owner or studio
- Studio distribution arm
- Content aggregator
- Television network
- Local television station

- Cable MSO
- Cable local unit

Some of the rules apply to VOD, some to SVOD, and some only to TOD. The key is that there are many rules that can come from any number of places. While it can seem daunting, it is quite easy to create and manage these rules.

Partitioning Metadata

The Video-on-Demand Content Specification as published by CableLabs has become the de-facto standard of how metadata is created and how it can incorporate many of the rules necessary to describe how on-demand content is to be handled. Initially written to support VOD (movies), it has been expanded to support SVOD. Moving forward, it is likely that the specification will need to be expanded to support all forms of on-demand content, including broadcast television.

Some metadata rules pertain to the specific content itself, while others apply to how that content is distributed and sold. One piece of content from a studio can be sent to many cable systems across the country. If the studio had to regenerate the content metadata each time, it would become a painful process that nobody would want to use. However, if the content specific metadata were attached or imbedded in the content itself, and the distribution specific metadata was separate, then the same content with metadata attached could be sent to many locations, with a different version of the distribution metadata. Thus, the content metadata and the rules-specific metadata has been partitioned.

1. Content Metadata

Content metadata includes program specific things such as a unique identifier, title, rating, description, time, actors, directors and crew, category, trailer file names, poster file names, etc. This type of metadata does not change, no matter who, what, when, or where it is distributed. This metadata could clearly be embedded in the actual content file and would stay with the file no matter where it goes.

2. Rules-specific Metadata

The rules-specific metadata starts at the content creation studio. The studio decides if there are any specific restrictions on the distribution and sale of this content and passes those rules along to the content distributors. For example, there may be a requirement to restrict a specific category of commercial - A “Friends” episode may require Coke commercials, but not Pepsi. From there, the studio distribution arm may require more specific rules. “Friends” may be allowed from Monday through Friday anytime, but not Thursday from 8-9 pm, to prevent intruding on first-run episodes. Further downstream, the television network may decide to allow viewing anytime on Tuesday and Wednesday because those are non-peak days. The local television station may want to restrict viewing from 10-11pm during the local news hour.

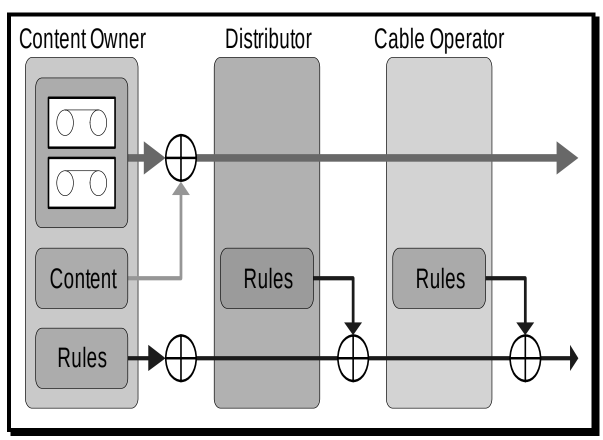

Figure 3-1 Rules-specific Metadata Flow

At each step along the way, the rules can become more restricted, but cannot be less restricted. In this manner, the content rules become more and more defined as they propagate downstream to the network operator and eventually the consumer (see Figure 3-1). Each system along the path is responsible for obeying the rules imposed upstream, and can expect each system downstream to obey the rules it passes on. When they reach the cable system, the TOD menu or EPG is built using these rules for the content received. By using this approach, the menus for the STB can be automatically and dynamically constructed.

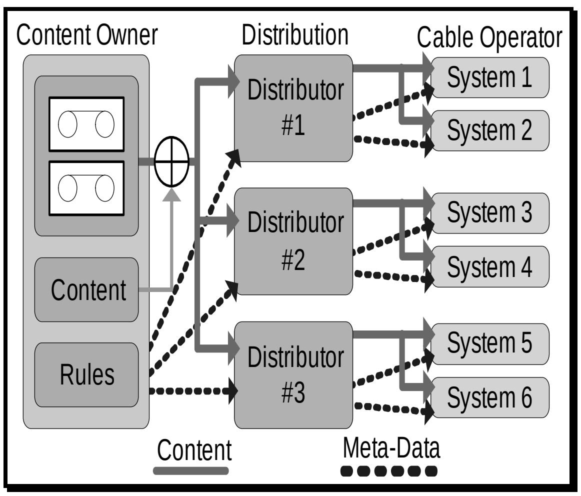

Figure 3-2 Metadata Flow to Multiple Downstream Paths

At each step in the process, there can be multiple downstream paths (see Figure 3-2) to both multiple distributors and cable systems. For example, the studio could sell a “Seinfeld” episode to the WB for certain nights in a specific week, and TBS on other nights. From each step facing down, the metadata can fragment, meaning there is a one-to-many relationship at each step of the way. This is important because at each level, a seller can sell to multiple customers. However, it would be inconvenient to have to re-record and re-master content each time it was sold. An improved solution would be to ship the exact same content to each downstream customer, but each would be supplied with unique rules-specific metadata which can be changed or updated at any time without requiring the entire piece of content to be resent.

Creating Metadata

With the two distinct types of metadata, appropriate software will be required to author and control its creation. A key ingredient is a unique identifier used to tie the asset together with both forms of metadata.

1. Content Metadata

The content specific metadata is created at the earliest possible point in the production and distribution chain. The best place for this is at the studio or encoding provider. In cases where the content is broadcast television, the content metadata could originate from the television network, or other production company supplying the network feed.

2. Rules-specific metadata

The rules-specific metadata can be created and adjusted at any point in the production and distribution chain, but would typically be originated at the same point the content is generated. For live television events, the rules could and should precede the actual content transmission. By sending the rules ahead first, the STB EPG can be populated, or other similar guide related decisions can be made.

Propagating Metadata

Both forms of metadata need to be sent along the same path as the actual content.

When any piece of content is sold or distributed downstream, the content metadata is included with the actual content along with an edited copy of the rules-specific metadata. Every copy of downstream content could have a unique set of rules-specific metadata, but the content metadata would stay the same. This allows each downstream provider to receive different rules, and allows them to be changed at a later time. When the rules change only the rules-specific metadata need be resent, not the content metadata or the entire program content. With this approach, any distributor in the chain can revise and update their rules-specific metadata as necessary.

Enforcing Rules-specific metadata

1. Asset Distribution

To make this system viable, each video server or file server along the asset distribution path must receive and obey rules encoded in the metadata. Typically in the role of asset distribution, all that is required is to pass-on the rules given to us. At any point in the path, the rules can be edited to become more restricted, but never less restricted. As assets are moved downstream to the cable plant, appropriate TOD software will pick-up the rules-specific metadata. The TOD software will use this rules-based data to build the availability matrix of programs, and associate a local time-slot for the consumer. The TOD server software is then responsible for ensuring that the studio/distribution/network rules and permissions are obeyed.

2. Content Propagation

When propagating content throughout the cable system, there can exist specific rules related to perishable content, or content that has a limited availability window. When this type of rule is implemented, it is important that the system remove such content and make the storage and streaming space available as quickly as possible. Another situation where the propagation of content is important is when a known high-concurrency program arrives and needs to be propagated to many places in a large network to facilitate the expected high demand.

CONCLUSION

In this paper, we have examined how conventional VOD servers are limited in their ability to ingest content and support the increasing stream requirements of TOD. There is a considerable impact in the output stream count as a VOD server is asked to ingest more content. With most existing systems, there is a non-linear loss of streaming capability while ingesting content. Specifically, many output streams may be lost for each single stream ingested. As the number of titles increases in VOD libraries the problem becomes more and more apparent. To reduce the impact on a VOD server, ingest of new content can occur after-hours. However this is just a temporary solution and won't scale as ingestion requirements continue to increase. With the upcoming everything on demand revolution, including Television on Demand, the ingest limitation of existing VOD server architectures becomes catastrophic. The more bandwidth consumed by ingest, the less bandwidth is available for streaming functions. Therefore more servers are required to keep the same stream count. As more servers are added, ingest and propagation becomes more and more complex. Elaborate ingest servers with content propagation services are a short-term solution but problematic longer term as unacceptable latencies are introduced to the distribution of broadcast television.

A new breed of servers designed specifically for Television on Demand is required. These servers need to handle over 100 streams of live ingest while simultaneously redistributing the ingested content to over 20,000 output streams. The server must not suffer any performance degradation in output streams while ingesting live or non-live content. The latency through such systems must be low enough to enable live television with trick-mode functionality similar to that of DVD. The streaming elements and the storage elements must be separately scalable and movable within the network.

With the plethora of ingested content from VOD, SVOD, and TOD, new means for authoring and propagating metadata must be implemented. In addition to content metadata, a new class of rules-based metadata will be required to protect revenue streams by allowing a rules-based distribution and STB presentation of content. The metadata must be partitioned and carried separately from the actual content to allow updating as well as customization depending on the MSO and region that the content is destined for.

A new breed of specialized, high performance TOD server with low-latency and live content ingest capabilities, plus a new metadata methodology, is a requirement to realize the potential of Television on Demand for cable operators.

About the Author

Robert Scheffler is Chief Architect at Broadbus Technologies, a provider of next-generation server systems that enable cable operators to effectively scale and migrate their networks from Video on Demand to Television on Demand™ – (TOD®)

Robert can be reached at:

(978) 264.7900
Robert.Scheffler@broadbus.com
http://www.broadbus.com