THE BENEFITS AND CHALLENGES OF DEPLOYING LARGE REGIONAL VOD ASSET LIBRARIES.

Michael W. Pasquinilli, Vice President of Engineering
Sunil Nakrani, Research Scientist
Jaya Devabhaktuni, Systems Architect
Concurrent Computer Corporation

Abstract

Many of the domestic video-on-demand (VOD) systems in service today are being upgraded to ten thousand hours of storage. Much of this storage is for traditional "on-demand" assets. There is also an industry trend towards recording an increasing amount of broadcast content onto the VOD system.

The cost of deploying these very large encode, ingest and storage libraries into each VOD system may prevent the launch of these new services. This paper discusses the economical benefits and technical challenges of introducing regional asset libraries that can support multiple VOD systems. The relationship between network bandwidth and asset caching will also be explored.

OVERVIEW

The current baseline storage level for VOD deployments is ten thousand hours for many of the major domestic MSOs. Ten thousand hours of content, including overhead, equates to approximately 24 terabytes of storage. Some of the more aggressive broadcast models have up to three hundred channels of broadcast programming recorded into the VOD system and retained for up to two weeks. If one hundred percent of this content were retained, the VOD system would require approximately 246 terabytes of storage. Even assuming that the cable operator is able to secure contractual rights for only twenty percent of this content, the storage requirements are still in the 50 terabyte range. (This does not take into consideration high definition (HD) content. These storage values could be three to four times larger for a system with all HD assets.)

Today there is a one-to-one relationship between the VOD system and the digital set-top box control system. The expansion of digital subscriber penetration is causing the digital set-top box systems to fragment into several mirrored digital video systems each with their own dedicated VOD system. This means that a single cable system today with one VOD system and one digital set-top box controller may soon split into three or four mirrored systems as digital subscriber penetration increases. With the current VOD architecture, each of these new digital video systems will require their own VOD library storage.

The economical challenge facing the MSOs of deploying very large VOD asset libraries in each digital set-top box system is further aggravated by the fragmentation of these single digital video systems into multiple mirrored systems. So whether the cable operator has a national footprint of many sites, or is a large single system operator, it is likely that with the current VOD architectures in place today duplication of the VOD asset library will be necessary.

To address this challenge several engineers in the cable industry are working towards developing a shared or "regional" VOD asset library that can serve multiple VOD systems. One of the major differences between the VOD library serving a local VOD system and a regional VOD library is that a local VOD library serves a closed network with dedicated network resources. In this case the local VOD library may "play" or "stream" the asset directly to the subscriber. However, in a regional VOD library, the network is likely shared with other data and is several Ethernet switches away from the subscriber. Issues such as QoS (quality of service), packet jitter, packet routing and trick mode latency make it less reasonable to expect a remote VOD library to stream across very large distances. In this case the regional VOD library must copy the asset to the local VOD system for play out.

The design concept discussed in this paper applies the hybrid VOD model in use today between the headend VOD library and hub VOD edge cache devices to a higher level in the VOD architecture. Now instead of a local headend VOD library, the VOD library is a regional library and, instead of hub VOD cache, the entire VOD system served by the regional library equates to the hub cache.

THE BENEFIT

There is clearly an economical benefit in capital savings if a single 100-300 terabyte VOD asset library could serve multiple VOD systems. Just the storage, ingest and streaming costs for such a system would be well over a million dollars. This does not take into account the costs in encoders, MPEG grooming, control systems, powering, cooling and operations.

The primary component that could undermine the economic benefit of a regional VOD library is the network cost to transport the video to each VOD system. This is where the hybrid VOD library architecture model comes into play. The primary benefit of a hybrid VOD library architecture is the savings in headend to hub transport costs. The challenge is to prove that this same savings could be applied to the regional VOD library.

THE COMPUTER MODEL

Since there are no regional VOD libraries in service in today's domestic cable market, it was not possible to gather measured data from an actual regional library. As an alternative, engineers at Concurrent developed a computer model that would simulate the operation of a regional VOD library. Real-world measured VOD asset usage data were used to exercise the model.

The Variables

The model had the following adjustable variables:

System Cache: The amount of memory available at the local VOD system to cache content pulled from the regional VOD library.

Time To Live (TTL): The amount of time that a downloaded library asset resides on the local VOD system cache before deletion.

Cache Management: A Least Recently Used (LRU) methodology was used to manage the local VOD system cache.

Network Bandwidth: This was the bandwidth assigned to each asset being pulled from the regional VOD library. For the sake of simplicity, a value of 3.75Mbps was used as a baseline minimum data rate for each asset pull. (It was assumed that all assets were MPEG 2 standard definition.) Multiples of this data rate were used to simulate a "best effort" data rate model. (It is more likely that a best effort methodology would be applied to this variable in an actual deployment.)

The model assumes that 100% of the VOD assets would initially be delivered to the regional VOD library. It was also assumed that the first copy of an asset delivered to the local VOD system would be cached locally. Assets would only be purged from the local system if the allocated cache value was exceeded. The TTL of an asset would be directly tied to the active usage of that asset by subscribers. If it was necessary to delete an asset from the local VOD system cache, the least recently used asset would be deleted to make room for the next requested VOD asset. No cached assets would be deleted that were in active use.

It was possible in the model to have denial of service. If the local VOD cache was full of assets being actively played by subscribers, additional requests for regional VOD library plays were denied. Since the objective of the model was to discover the network bandwidth necessary to support the regional VOD library, no restrictions were placed on the network bandwidth between the regional VOD library and each of the local VOD systems.

In several runs of the model the value of the local VOD system cache was changed. We were looking for the amount of local VOD system cache that would allow for no denial of service and have the largest impact on reducing the network bandwidth between the regional VOD library and the local VOD system.

For the purpose of this simulation, actual VOD asset usage data was collected from three large regionally co-located systems that have network connectivity via the MSO's internet backbone. Figure 1 shows a simplistic block diagram of the configuration.

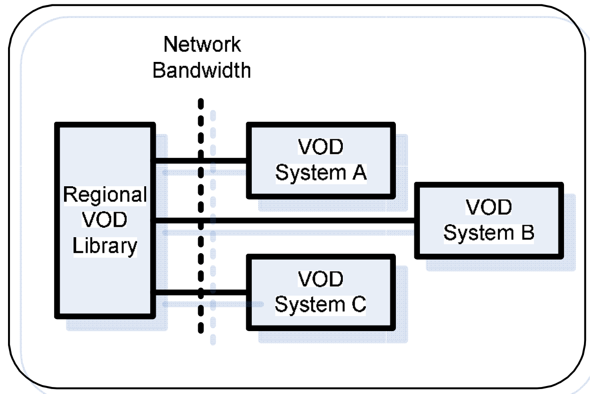

Figure 1. "Three VOD System Model"

The Results

The first run of the model assumed that each individual asset transfer rate from the regional library to the local cache was fixed at 3.75Mbps. No cap was placed on the network bandwidth between the regional library and the local VOD system. The local VOD system cache value started at 250GB and was incremented by 250GB until no Denial of Service (DoS) was encountered due to lack of local VOD system cache storage. Results of the first run are shown in Table 1 below.

System	Cache @ 0% DoS	Peak BW	Avg BW
A	2.25TB	1.67Gbps	0.61Gbps
B	3.0TB	2.47Gbps	0.83Gbps
C	1.0TB	0.90Gbps	0.32Gbps

Table 1. "First Run At 3.75Mbps/Asset"

In the next run all other variables were kept constant, but the asset transfer rate was increased to 15Mbps per asset. Results of the second run are shown in Table 2 below.

System	Cache @ 0% DoS	Peak BW	Avg BW
A	2.25TB	3.7Gbps	0.645Gbps
B	3.0TB	4.4Gbps	0.870Gbps
C	1.0TB	1.4Gbps	0.335Gbps

Table 2. "Second Run At 15Mbps/Asset"

Analysis

Although System C had the highest number of VOD streams during the data collection cycle, and was very near the stream value of System B, it required the least amount of local VOD system cache to store assets pulled from the regional library. This was due to the following: In System B the top 20% of assets accounted for 83% of the views. In System C the top 20% of assets accounted for 88.6% of the views.

This shows that the usage patterns of the subscribers in the system can have a direct and measurable impact on both the local cache and network bandwidth necessary to support a regional VOD library. This is shown to be true even when the two systems have identical content offerings and very similar overall stream usage.

Another important item to note from this data is that a four times increase of the individual asset data transfer rate resulted in a negligible increase in the average data rate and a less than double increase in the peak data rate. This points to the benefits of allowing a best effort transfer of the library assets up to the ingest data rate of the local VOD system. Using the local caching method resulted in a combined peak network bandwidth of only 9.41Gbps. If no local caching were used and instead each asset was streamed directly from the regional library, the combined peak network bandwidth required to support these three VOD systems would be 102Gbps.

TECHNICAL CHALLENGES

The technical challenges facing us today are not in fielding these large VOD asset libraries. We have demonstrated the ability to support large storage systems with high ingest rates and very low encoded content throughput latency. We can see from the data above that both the network bandwidth and cache storage values are manageable. The primary technical challenge facing the deployment of regional VOD libraries is in the fact that the VOD standards and architectures in place today were not designed to accommodate content not entirely under the control of the local VOD back office controller.

The following section will outline some of the major areas that are being addressed in order to support a regional VOD asset library.

Metadata Publishing

The most common way that an asset is received into a VOD system is via a satellite "catcher." These catchers pass over to the VOD system the asset and the metadata file. The VOD back office manages placement of the metadata into the back office database and the transfer of the asset into the VOD server(s). There are back office checks in place today to make sure that the metadata and assets are both accounted for in the VOD system.

With the regional library assets, the ingest point for the asset is at the regional library. Therefore there must be another entity outside of the local back office that is keeping track of the metadata for the assets stored on the regional library. This *other* system must then publish to the local back office the metadata of the assets stored in the regional library. This needs to be done so that the local back office can manage the provisioning, rules and lifecycle of the asset once it is within local control. Most all of the VOD set-top box navigators in use today pull their data from the local back office. This being the case, the local back office must have all the metadata for both local assets and regional library assets within its database.

The problem here is that the local back office does not currently have the means to accept and act on metadata whose associated assets are not locally ingested. (The Comcast NGOD set of specifications begins to address many of these issues, but as a proprietary specification cannot be discussed in this paper.) A new method needs to be adopted that allows for the regional library asset metadata to be identified as representing remote content and changes must be made to the back office to recognize and manage assets not stored locally.

Trick Modes

It is most often the case that VOD trick mode files and/or indexes are created during asset ingest. This trick mode creation usually happens as the asset file moves from the catcher into the VOD system. Each VOD pump vendor has their own methods of creating trick mode files/indexes.

Since it is envisioned that a regional VOD library would support multiple VOD pump vendors, the trick mode creation would not happen at the ingest point at the regional library, but must instead happen at the time when the asset is delivered from the library to the local VOD system. Depending on how this process is done, it could result in commercially unacceptable latency for the subscriber. (An alternative approach considered is requiring all VOD vendors to adopt a common trick mode standard.)

A second consideration related to trick modes is the effect of the fast forward function. It could be possible that a subscriber tries to *fast forward* to the end of the asset faster than the file transfer speed for that asset. This is why I mentioned earlier in this paper that a best effort methodology of file transfer from the library to the local VOD system would be preferred over the fixed "play rate" of the title. Otherwise the regional library must somehow support the function of trick mode play out especially for the fast forward function.

Latency

An ideal architecture would have the subscriber oblivious to the fact that a VOD title was being served from a regional library versus a local VOD server. But this behavior is not currently guaranteed.

With proper network provisioning and QoS, the path from the regional library to the local VOD system should not contribute to the latency. However, the methods for how the local VOD system ingests and creates trick modes and propagate content will be the most likely cause of latency. Some VOD systems may not be able to play an asset until the entire asset is copied into the VOD server. This restriction may also apply to the creation of trick modes. Depending on the ingestion point into the system and the methods of content propagation, there could be queuing delays before the asset arrives at the VOD server designated for play out.

Asset Lifecycle

Today the lifecycle of a VOD asset is defined by the metadata associated with the asset. The VOD system will retain the asset for as long as is specified by the metadata. Since both the regional library and the local system would have access to the metadata, both systems can continue to use this information. However, for the purpose of managing local storage, a library asset "copy" must be marked eligible for deletion prior to the asset expiration date.

If a local system allocated one thousand hours of storage for caching library content, some content would need to be deleted at times to make room for other requested library content. The regional library may contain many tens of thousands of hours of content and not all content pulled down to the local site could necessarily be retained for the full metadata-specified viewing period. Therefore assets pulled from the regional library must be "allowed" to be deleted to adequately manage the local cache.

Bookmarks, Active Rentals, Resume Viewing

One of the major challenges of a regional library is supporting the subscriber experience of being able to view an asset multiple times within the rental window (usually 24 hours). To understand the issue consider the following scenario: The subscriber watches all but the last ten minutes of an obscure video that was pulled down from the regional library. The subscriber returns some twenty hours later to resume viewing the remainder of his bookmarked video. Since this was an obscure video it was likely purged from the local cache to make room for more popular or recently requested assets. In this case the subscriber only wants to watch the last twenty minutes of the asset. Does the library download the entire asset again to the local cache or a partial file? Since it may be too difficult to push bookmark information all the way back to the regional library, and it may be too difficult to manage file fragments within the local cache, it is likely that the entire asset must be copied again to the local system.

Ad Zones

The largest regional VOD libraries are likely to be those made up of broadcast content supporting "Look Back" and other network DVR types of services. Ideally the cable operator would like to record just one copy of a regional broadcast channel. Ad zones make this difficult. In our three system examples, each of these systems may have had four or more local advertisement insertion zones. One method of dealing with ad zones is to record one copy of a broadcast for each zone into the VOD system. These broadcast recordings would start to add up as they are multiplied by the number of channels that have local advertisement opportunities, times the number of ad zones, times the number of systems supported by the regional VOD library.

An alternative way of dealing with ad zones is to insert the local advertisement at the point of play out. In our case this would be at the local VOD system's VOD server. Since VOD provides a dedicated session to each subscriber on demand, there is an opportunity to target ads at the subscriber level or to at the very least keep the ad zones intact.

CONCLUSION

Based on the computer model and the measured VOD asset usage data from just three VOD systems it would seem that a LRU managed local VOD system cache of about 1,000 hours per system and total network bandwidth of at least 5Gbps (peak) per system would be a good starting point to support a regional VOD asset library. This does not take into consideration the impact of high definition content on these variables. Also, this specific model does not include "Start Over" type content usage.

It is important to reiterate that a change in subscriber usage patterns can significantly change the resources necessary to support a regional library. This being the case, computer modeling that utilizes actual measured asset usage results will be more important in defining resource requirements than will anecdotal or historical experiences.

There are quite a few technical challenges to be addressed before a regional VOD library is commercially viable and transparent to the subscriber. These challenges represent changes to the VOD back office, VOD system and VOD server. When all of these challenges have been met the VOD system will have evolved into a video delivery platform that will satisfy all of the on-demand and broadcast needs of the MSO at the national, regional and local level.