Using the Cloud and Edge Video Servers as an Alternative to Satellite for the Distribution of Linear Channels

Pablo Salomon, CEO Inetsat
Bruno Martinez, CTO Inetsat

Abstract – *We propose a "**distributed playout**" distribution model for linear channels with no live content that has many advantages in terms of cost and flexibility compared with the traditional satellite distribution model. In this model, satellite Integrated Receiver/Decoders (IRDs) at the local operator headends are replaced with low cost video servers connected to low bandwidth internet connections to a backend platform in the cloud. Each video server stores enough content so that previously aired content does not need to be distributed again in subsequent repetitions and the content and playlist are distributed several hours before air time to ensure continuous playback even if the internet connection goes down for a few hours. The cloud backend monitors and updates the video servers with new content files and playlist, allowing the automated operation of hundreds of video server to be as simple as the operation of a single server. This same architecture can - also be used in combination with satellite for localization/ad insertion purposes.*

Traditional Distribution Models

<u>Satellite distribution model</u>

Most pay TV channels use satellite distribution to get their feeds to operators at different locations. A playout center generates a real-time feed on a central location based on live content and pre-recorded content files. This signal is transmitted into a satellite in encrypted, compressed form, and the satellite in turn re transmits the signal to its area of coverage (footprint). The cable operators receive the stream from the satellite through an integrated receiver decoder typically supplied by the channel, multiplex all channels together, and distribute them to the audience through their local network typically using fiber and/or coaxial transport.

<u>Point to Point distribution model</u>

Another model sometimes used as an alternative to satellite is having direct point to point terrestrial IP links with each local operator typically using fiber. As with the satellite distribution model, this model also requires a playout center, operated by the programmer, where the feed is generated in real-time based on live content and pre-recorded content files.

Using internet connections instead of direct point to point links is usually problematic as bandwidth needs to be guaranteed in order to maintain broadcast quality and the connection availability is critical and any interruptions in the connection will cause the channel to be off the air.

Proposed Distribution Model

<u>Distributed Playout Model</u>

This model proposes to create the feed playout simultaneously on multiple low cost video servers at the operator's headend. These video servers should be connected to the internet to low bandwidth broadband internet connections. For most channels 1 to 2Mbps connections should be sufficient.

Each video server stores enough content so that previously aired content does not need to be distributed again in subsequent repetitions and the content and playlist are distributed several hours before air time to ensure continuous playback even if the internet connection goes down for a few hours.

Through the internet connection each video server frequently polls a backend system hosted in the cloud which monitors and updates the video servers with new content files and playlist.

This backend should automate the synchronization of all the video servers associated with the same feed in order for the operation of hundreds of video server to be as simple as operating a single server.

<u>Advantages</u>

The key advantage of this model is its ability to replace the high cost satellite or fiber point to point connections that must have guaranteed bandwidth with low cost, low bandwidth internet connections with no need for guaranteed bandwidth or quality of service while achieving higher reliability and always high image quality on the feed output to the operator. In order to achieve this key advantage, channels must have some characteristics described in the following section.

Another advantage is not requiring any equipment on the programmer or broadcaster side as no central playout or satellite uplink is needed. The content can be uploaded to the cloud and the channel can be managed from any desktop PC. Also, the internet connection required has the same low bandwidth requirements as the connection at the operator's side. Even if the content needs to be distributed to hundreds of video servers, the channel will only need the bandwidth to upload it once to the cloud and then all the video servers will download it from the cloud.

<u>Channels that could benefit from this model</u>

In order to benefit from this model, channels must have the following 2 characteristics:

1) Repetition: As the proposed model sends the content only once to the video servers on the operator's headend where it is stored and reused on all its repetitions, the higher level of repetition means the lower amount of unique content that must be distributed and therefore the lower bandwidth requirements.
2) No live content: This allows the content to be distributed ahead of air time and for each hour of content to be transferred over multiple hours. This enables the distribution bandwidth to be lower than the content encoded bitrate and it doesn't require the distribution bandwidth to be constant.
Except for sports and news, most pay TV channels today meet these 2 characteristics. Most of the content is available weeks in advance. It's not unusual for content to be repeated in 6 or 8 hour blocks and to be repeated again in the following days and months. For example, a movie channel typically has between 50 and 100 hours of premieres a month of the 720 hours available each month.

OPERATIONAL CONSIDERATIONS

Transparency from the Operator's perspective

For the proposed solution to be successfully deployed, the video servers should be as similar as possible to IRDs and avoid introducing any overhead or new practices for the operators. This includes providing video servers of the same size as IRDs so they can be racked together, having the same outputs as IRDs so they can be connected to their existing distribution platform and having fully automated operation so that no management is required at the operator's headends.

Optimizing the use of bandwidth and storage resources

In order to minimize the bandwidth requirements and the storage cost both at the cloud and at each video server, media files should be transcoded before they are distributed to the quality required by the operators. In other words, if operators are going to output an Mpeg4 stream at no more than 10Mbps from the video servers, it is a waste of resources to distribute files encoded at a much higher bitrate.
Although the video servers should output the transport stream at constant bitrate to meet the requirements of most operators, media files could be encoded at variable bitrate for distribution to further optimize bandwidth for a specific level of image quality.

Optimizing the use of processing capacity at the video servers

One possible way to minimize the use of processing capacity at the video servers would be to distribute the media files in transport stream format so that the video server will just need to read the bits from the files and send them through the output with minimal or no processing. However, for most channels this would not be an option as they typically need to overlay different secondary events on every repetition of the content. Therefore the video servers will require the capacity to decode the video of the media files, overlay the secondary events as needed and encode the output stream. As the audio is not typically affected by secondary events, the pass-through approach that requires no decoding and encoding may be used for audio channels if they are encoded on the media files with the same specifications needed at the output stream.
On the video, other types of processing may be avoided if the media files are encoded at the same resolution, frame rate and chroma subsampling as the ones needed at the output stream.

Bandwidth Requirements

The following formula can be used to calculate the minimum bandwidth required for the content distribution. This is the downlink bandwidth at each operator's headend and uplink bandwidth at the channel side.

$$B = \frac{E*H}{720}$$

Where:
B is the bandwidth required
E is the bandwidth at which the media files are encoded
H is the amount of hours of new releases a month that where not previously distributed to the video servers.
720 is the amount of hours in a 30 day month.

This is a theoretical formula assuming a perfect internet connection with constant bandwidth and that content is available with enough lead time so that there is always content being distributed. In a real life scenario, the internet connection should have additional bandwidth as a margin to contemplate not so ideal conditions.

Content Distribution

Content should be distributed to the video servers based on when they will be aired. This coordination requires joining the playlist data with the records of which video servers already downloaded which content files and prioritize the download accordingly. Content that has already been uploaded to the cloud but is not yet scheduled on the playlist may still be distributed to the video servers with a lower priority.
There are more complex scenarios in which a channel may use this model to distribute multiple localized versions of the channel where part of the content is going to be used by all the video servers while other content may be specific to a few video servers in a specific region. This requires a more complex media asset management in the cloud in order to avoid distributing unnecessary content to video servers that will not need it wasting both bandwidth and storage.
Another case is when a new video server is added to an existing deployment. While the new content grows slowly, an empty video server will need to catch up with all the content already distributed. A possible solution is to preload that video server before being sent to the operator's headend

or to connect it to the internet weeks before it is set to start broadcasting so that it can catch up with all the previously distributed content as well as the new content that is being distributed for future broadcast.

HOW IT COMPARES?

In this chapter we compare the proposed model of distributed playout with the traditional satellite distribution model across several factors.

Equipment Cost (Capex)
In the traditional satellite distribution model, the equipment required to distribute a new channel consists of:
a) At each operator's headend: An IRD and a Satellite dish which may be shared among multiple feeds.
b) At the broadcaster/programmer: A complete playout chain from the playout server where the feed is originated up to the uplink satellite dish. As this playout chain is a critical point of failure there is usually additional costs involved in having redundant equipment.
In the proposed model:
a) At each operator headend: A video server is required to playout the feed.
b) At the broadcaster: No special equipment is required. The channel can be managed with a standard desktop PC.

As described above, the only Capex cost in the proposed model is the cost of the video servers at each operator's headend. As of 2013 it is possible to build such video servers using off the shelf components for a total hardware cost of under USD 2000. This is more expensive than an SD IRD or a low end HD IRD but cheaper than high end HD IRDs with transcoding capabilities like the ones provided by a video server. This comparison does not take into account the software licensing cost for the video servers which vary significantly.

Distribution Cost (Opex)
While pay TV channel revenues tend to be proportional to their reach, in a satellite distribution model the satellite transponder cost and additional playout and uplink services are fixed for a specific feed & bandwidth regardless of the amount of operators receiving that feed.

On the proposed model, the operational costs are:
a) The cloud storage: This is the only cost not proportional to the amount of operators distributing the channel. For example, at Amazon S3 prices as of January 2013 of USD 0.08 per GB per month, having 2TB (aprox 1000 hours of SD content) stored would cost USD 160 per month.
b) The cloud bandwidth for content distribution:
This cost is proportional to the amount of headends that the content would need to be distributed to. At Amazon S3 Data Transfer prices of USD 0.12 per GB, distributing 200GB of media files (aprox 100 hours of new content) would cost USD 24 per headend per month.
c) The internet connection at each operator's headend: Most headends would already have an internet connection that they could reuse for this purpose without any additional expenses. However if it is necessary to have a separate internet connection, most low end broadband connections should be sufficient. Depending on the location, these connections typically cost in the range of USD 30 to 50 per month.
d) The licensing cost of a software platform used: Both the pricing and the business model will vary depending on the provider chosen. For example, providers like Inetsat offer a model proportional to the amount of headends.

Adding all the operational costs, for feeds that need to be distributed to a low amount of headends (under 10), the proposed model can be up to 10 times cheaper than the traditional satellite distribution model. Even for 100+ headends, these operational costs are significantly lower than the equivalent satellite distribution costs.

This evaluation does not take into account the personnel cost to manage the distribution which should be equal or less the proposed model vs. the traditional satellite model.

Time Shifting Cost
When the same feed serves locations in different time zones, some viewers may not get programs at the desired time.
To solve this, additional time shifted feeds may be distributed.
In the traditional satellite distribution model, each additional feed will have about the same cost as the original one, effectively multiplying the distribution cost.
On the proposed model, there are no underlying cost changes as the same playlist and the same content is distributed to the time shifted video servers. All that is needed is to adjust the time zone of each video server to the desired time zone for that specific operator.

Localization Cost
As with the time shifting costs described above, having multiple localized feeds effectively multiplies the distribution cost in the traditional satellite distribution model.
In the proposed model, none of the underlying costs change by having multiple localized feeds. The only additional cost may be licensing cost of the software used if its provider takes into account the amount of feeds on the pricing.

Reliability
In the traditional satellite model, there are multiple critical points of failure along the distribution chain that would cause the channel to go off air in all operators. This includes the complete playout chain at the facility where the channel is originated as well as the satellite itself. To mitigate this, it is common to have multiple redundant equipments at the facility where the channel is originated however it is much

less common as it is very costly to have multiple separate facilities and multiple satellites for redundancy.
The architecture in the proposed model has no single point of global failure. This architecture gracefully handles the failure of the internet connections or even if the cloud infrastructure is unavailable, as the video servers continue to function relying on the already distributed media and playlist. A video server can fail, causing the unavailability of the channel at a specific operator but this can be mitigated via simple redundancy (having 2 video servers at each operator's headend).

Image Quality
In the satellite distribution model there are several factors like extreme weather or solar storms that can interfere with the satellite transmission affecting the image quality of the feed.
With the proposed model, a perfect quality of the feed is always delivered at each operator's headend.

Global Coverage
Satellites have a limited geographical coverage. In the traditional satellite distribution model, in order to deliver a channel into multiple continents, multiple satellite transponders must be used as well as multiple facilities to perform the uplink of the feeds. This multiplies the distribution cost.
By using the internet as the transport for the content distribution, the proposed model has global coverage and can serve operators on multiple continents without any additional costs.

HYBRID MODEL

As described in the "Channels that could benefit from this model" section, the main limitation of the proposed distribution model is that it cannot be used for channels with live content.
However, channels with live content could adopt a hybrid model to gain the benefits in terms of localization of the proposed solution while reducing their distribution costs.
This hybrid model consists of using the traditional satellite model to distribute live content and prerecorded content that needs to be distributed across all the operators while using the distributed playout model to distribute the localized content specific for each operator or region.
At each operator's headend, there will be both an IRD and a video server. The output of the IRD will be connected to an input on the video server which will perform the splicing and pass through the satellite feed or play out the localized content as needed. A signaling mechanism on the satellite feed like SCTE 35 or cue tones should be used to synchronize the timing and perform the transition with precision.
For example, a channel that has 5 localized feeds where part of its content (including all the live content) is shared among them could use a single satellite feed to distribute that shared content and perform the localization of each feed using the distributed playout model while significantly reducing their total distribution cost.
This hybrid model is similar to ad insertion solutions already available in the market.

CONCLUSIONS

Many channels could benefit from the distributedplayout model proposed on this paper.
The three main factors that enable this model to be deployed today while a few years ago it was difficult and not cost effective to deploy are:
1) The availability of low cost broadband internet connections in almost every city around the world.
2) The availability of cloud computing platforms that makes it possible to store and distribute large amounts of data at low cost and makes it easy to scale.
3) The availability of powerful processors, multi-terabyte hard drives and ASI/SDI cards at low prices offer the possibility to build a video server with off the shelf components at a cost similar to an IRD.

As this model can lower the distribution costs in some cases up to 10 times, it has the potential to disrupt the traditional satellite distribution model used today while enabling new niche channels or localized feeds that were not commercially feasible due to the high fixed cost of satellite distribution.

AUTHOR INFORMATION

Pablo Salomon, CEO Inetsat – psalomon@inetsat.com
Bruno Martinez, CTO Inetsat – bmartinez@inetsat.com