Transcoding Choices for a Multiscreen World

Jonathan Ruff, Global Technical Marketing Lead

Santhana Chari, VP Engineering Digital Video Systems

Mark Depietro, VP Corporate Strategy & Business Development

Introduction

The multiscreen video movement keeps growing steadily. As household penetration rates surge for tablets, PCs, gaming consoles, smartphones, connected TVs, and other video devices, growing numbers of cable TV operators, telcos, satellite TV providers, and other pay TV distributors are jumping into the market with offerings that deliver video to the full range of consumer screens. As these video service providers increasingly seek to deliver their video offerings to multiple devices, one of the greatest challenges they face is how to deliver high-quality content in the right format, securely-without overtaxing the access network.

Using adaptive bit-rate streaming, providers are now able to transcode their content for virtually any device. But as the number and type of devices keep multiplying, the process is growing increasingly difficult to manage. In this white paper, we will explore the applications for home- and network-based transcoding, and preview some of the innovations that are emerging to help providers transcode their content more efficiently and effectively in the multiscreen world.

Network vs. Home-Based Transcoding

There are two basic architectural options that service providers can choose for video transcoding. The first is to carry out the video transcoding somewhere in the provider's video distribution network. This can either be near the core of the provider's network, which has benefits for national programming, or closer to the edge, which streamlines transcoding for regional content. In either case, this approach centralizes the transcoding resources and associated costs, enabling greater efficiencies as take rates improve.

The second basic architectural option is in-home transcoding, where the provider relies on a headed or headless gateway in the home to transcode the incoming video signals for playback on devices beyond the set top box. With this approach, the same video content, in the form of legacy QAM delivery, is sent to every home and then transcoded for the consumer's device in real time. While this approach uses less network bandwidth, it requires the expense of a transcoder in every subscriber's home.

While some video solutions vendors are steering service providers to either network- or home-based solutions, transcoding is not a "one size fits all" proposition. Several factors can help service providers determine which is the right approach, including the type of content being distributed, the number of devices supported, the applications and associated protocols used, and the capabilities of the access network. There may even be cases where it is advantageous for service providers to deploy both network- and home-based transcoding. It all comes down to choosing the right approach based on the service provider's service goals and system capabilities.

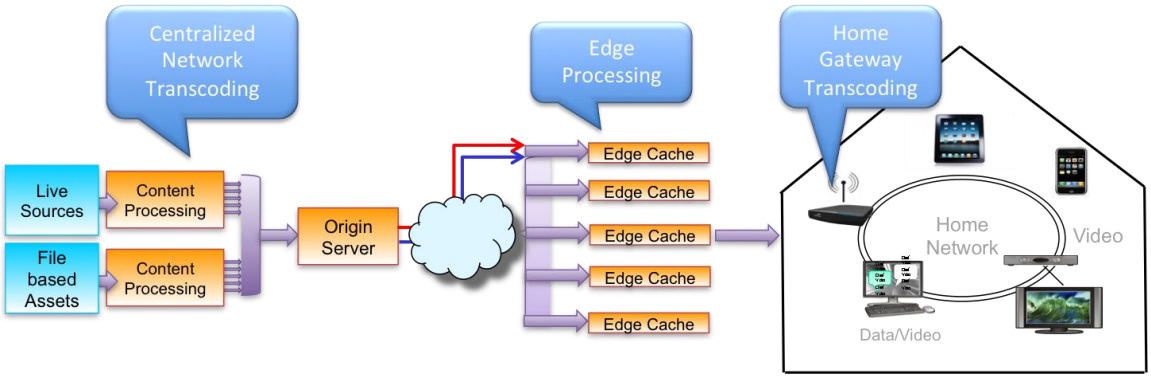

Figure 1: Network and Home-Based Transcoding Locations

Choose Which Content to Stream

As it is with so many video service provider decisions, the content being offered to video subscribers vastly impacts whether home- or network-based transcoding makes the most sense. In a multiscreen world, there is a sizable difference between live and stored content, and subtle differences in the balance between the two can have dramatic effects on the decision to transcode in the network versus in the home.

In the case of live TV, streams must be transcoded in real time and ingested into the CDN, and the more channels a service provider wants to offer beyond the set-top, the more complex the network-based transcoding requirement becomes. With network transcoding, each client request is delivered as a unicast stream over the access network. As the number of linear channels offered and take rates increase, these unicast streams can begin to consume significant bandwidth resources, making home-based transcoding a more scalable solution for delivering linear broadcasts.

Just as with live TV, stored content carries its own set of transcoding challenges. This time, the decision to transcode in the network versus the home depends on where the content is stored and where it will be made available. For centralized content such as network-based DVR and on-demand programming, a network-based transcoding approach can offer significant bandwidth savings by delivering only the content profile needed for the viewing device being used. In addition, by transcoding in the network the content can be made available ubiquitously to subscribers, no matter where they are.

Home-based DVR content that is being viewed in the home is far better suited to a home-based transcoding approach, since content would otherwise need to be transmitted up through the network for transcoding and back to the home for viewing, which is subject to the available bandwidth in the upstream network and is impractical

when congestion is high. In addition, home-based transcoding becomes problematic when the same content needs to be viewed outside the home because of today's upstream bandwidth limitations. As an increasing number of users seek to take their content with them, shared upstream network resources can quickly become exhausted. In this scenario, a "sync and go" approach may be the most practical solution.

Support the Right Mix of Profiles

One of the greatest challenges for video service providers today lies in keeping pace with the myriad devices consumers are using to access multiscreen video. Consider the number of display resolutions, streaming protocols and conditional access methods that need to be supported in order to deliver multiscreen video to every user. Table 1 represents a small sampling of these variables, which, when combined, require support for over 100 unique video profiles.

Maximum Display Resolution	Streaming Protocol	Conditional Access Method
480 x 320 (HTC Desire C)	Apple HLS	PlayReady
800 x 480 (Samsung Galaxy W)		
960 x 640 (Apple iPhone 4s)	Adobe HDS	
1024 x 600 (Kindle Fire)		SecureMedia
1024 x 768 (LG Intuition)	Microsoft Smooth Streaming	
1136 x 640 (Apple iPhone 5s)		
1280 x 800 (Motorola Xoom)	MPEG DASH	Adobe Access
1366 x 768 (Apple Macbook Air)		
2048 x 1536 (Apple iPad 3)		

Table 1: The Many Video Profile Variables

For service providers seeking to deliver multiscreen video, this profile explosion presents a dilemma. They must choose between creating an environment where the number of variables can be kept to a minimum, or dedicate significant resources to supporting every profile under the sun. Each of these approaches comes with trade-offs, and each effects the decision of transcoding in the network versus the home.

There are multiple ways to minimize the number of profiles that need to be supported. One is to allow consumers to access multiscreen services only on the newest devices running the latest software. Alternatively, service providers may choose to limit their multiscreen offerings to devices made by certain manufacturers or running certain operating systems. In either case, the cut-off point should reflect the popularity of devices that consumers are using to access video content, where the supported smart phones, tablets and laptops encompass the majority of the market. By limiting the total number of devices that need to be supported, service providers can drastically reduce the number of profiles that need to be created and streamed, making network-based transcoding a practical solution.

By instead committing to support every combination of resolution, streaming protocol and conditional access method, service providers are embarking on a journey that is well suited to just-in time transcoding, which can be performed in the network or at the home. Todays' powerful transcoders are adept at processing the multitude of variables and creating the right profile for virtually any device in real time. While performing real-time transcoding in the home would require each gateway to be updated frequently in support of new device requirements, doing so in the network would consume significant bandwidth resources.

Evaluate Network Capacity Requirements

A key factor in the decision to transcode in the network versus the home is the capacity available in the access network for multiscreen video. Generally, transcoding in the home is less burdensome on the access network, since each linear channel or stored content file can be delivered to the home only once. Since service providers rarely have bandwidth to spare, it may be useful to determine how much capacity is needed before deciding where to transcode.

Doing so often means conducting usage modeling that takes into account the anticipated multiscreen traffic requirements at various times of day. These network dynamics are changing as consumers adopt the multiscreen lifestyle. For example, based upon a recent study conducted by Motorola Mobility (which is now part of Arris), the multiscreen world now has its own 'prime time' slot - the late evening hours when viewers are increasingly bringing tablets into the bedroom to watch TV.

While usage modeling will help determine how much capacity is needed to stream multiscreen content, service providers will need to create efficiencies wherever possible to accommodate this increase in network traffic. This means supporting multiple bit rates to accommodate network congestion, utilizing new compression technologies and applying intelligent transcoding practices.

Since multiscreen content is delivered over an IP network where network congestion is inevitable, service providers have found a solution in creating multiple bit rate versions for each profile. This ensures that the highest quality content is delivered when network capacity permits, while sending a lower quality version during times of congestion. However, this adds yet another variable to the number of profiles that must be created and sent. For the example in Table 1, supporting three bit rates for each profile now requires the creation of over 300 distinct versions of each piece of content. While all of these versions will not have to traverse the network at a given time, the capacity requirements for the network-based transcoding approach scale linearly with the number of client requests and therefore can still be significant.

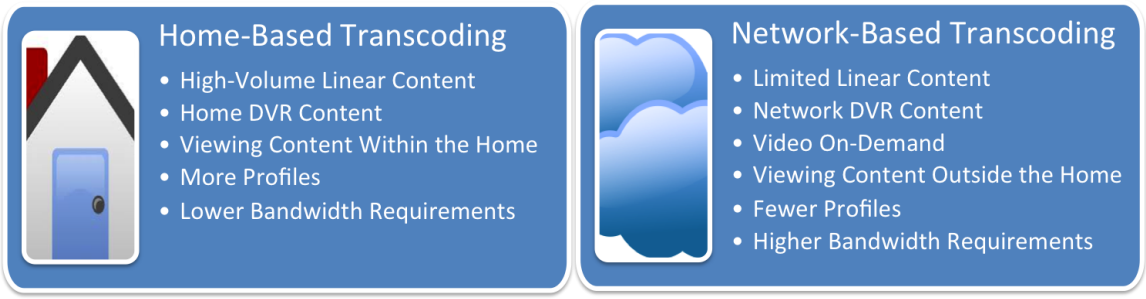

Figure 2: A Summary of Home- and Network-Based Transcoding Dynamics

Fortunately for service providers seeking to leverage network-based transcoding without overtaxing the network, a new compression standard is emerging as a potential game changer. While HEVC is being developed primarily as an enabler of Ultra High Definition Television (UHDTV), its earliest impacts may be felt not on the television, but the other screens in the house. HEVC, the successor to MPEG4, cuts the bandwidth requirement for video in half - without reducing picture quality. This creates the opportunity for service providers to continue supporting three bit rates for each profile without breaking the bandwidth budget. HEVC can be implemented in software on client devices, making it an unobtrusive upgrade that may soon be credited as a key enabler of ubiquitous network-based transcoding.

Prepare for Future Multiscreen Innovation

Beyond supporting multiple bit rates and utilizing HEVC for significant bandwidth savings, service providers are investigating several new innovations that may further clarify the transcoding decision. With new technologies emerging that streamline content delivery outside the home, reduce the storage requirements for encoded content, and intelligently automate the profile selection process, both home- and network-based transcoding can be made more efficient, practical and affordable.

As discussed above, one of the key challenges to home-based transcoding lies in the ability for subscribers to access content when they are outside the home. While

"broadcasting" from the home may not be practical on a widespread basis due to upstream bandwidth limitations, adding intelligence to the sync-and-go process may provide a viable solution. Rather than asking users to decide which content they would like to load onto their devices before leaving the home, intelligent applications might instead communicate with the home gateway and automatically sync the content that is most likely to be watched. The application may choose content that is most recent, or best matches the user's stated preferences or viewing activity. This *intelligent sync-and-go* solution has the potential to overcome one of the greatest challenges to home-based transcoding- upstream bandwidth usage.

Just as intelligent sync-and-go can help overcome network limitations for home-based transcoding, *just-in-time packaging* has the potential to ease the processing and storage burdens for network-based transcoding implementations. Just-in-time packaging calls for transcoding of the content for all profiles required, but only packaging or "chunking" of the content as subscribers request it. This significantly reduces the number of versions of each piece of content that need to be saved, and eliminates the creation of versions that are never required. By driving more efficiency into the process, just-in-time packaging has the potential to make network-based transcoding an even more attractive approach.

In many cases, today's transcoders are creating versions of content based on static profiles, and depending on the nature of the content, some of the versions are not always required. By better understanding the dynamics of the content, transcoders may soon be able to evaluate content before fully processing it, and eliminate creating versions that are unnecessary. This *dynamic profile selection* capability may soon emerge as a means of further reducing the number of versions of content that need to be created, stored, managed and streamed in network-based transcoding implementations.

Conclusion

Given the multitude of variables involved, including the content that needs to be supported, the profiles that need to be managed and the unique network dynamics at play, deciding between network- and home-based transcoding is anything but an easy choice. And in a period of increasing transcoding innovation, the choice doesn't seem to be getting any easier. But for service providers, the good news may just be that – at least in the near-term – there is no single correct answer.

The decision to transcode in the home versus the network is vastly impacted by each service provider's unique service goals, customer requirements and network conditions. Just as home-based transcoding is best suited for whole home DVR implementations, network-based transcoding wins the day for network DVR deployments. While transcoding in the network is best for seamless delivery of content in and outside the

home, transcoding in the home is best to minimize access network bandwidth requirements. And even if one approach seems perfect given the content and device requirements – network limitations may come into play and change the game completely.

Choosing the right transcoding approach may not come down to an either/or decision at all. Instead, it is more likely that the correct and chosen method will vary based on the service provider, network segment, or even the service being offered. Today's debate on home- versus network-based transcoding may even go the way of to yesterday's battle between home DVR and video on-demand.

Where industry pundits were once polarized in favor of either home versus network storage for *time*-shifted content, we are now equally divided over home versus network-based transcoding of *place*-shifted content. And once again, today's correct answer might just be a combination of the two. But as more access network bandwidth becomes available through node splits and DOCSIS 3.1, and new efficiencies are unlocked using HEVC, we may one day see network-based transcoding win the day. Wherever the debate ends up, the key for service providers is to develop an intelligent transcoding strategy that meets their unique needs today as their subscribers increasingly move into the multiscreen world of tomorrow.

Related Reading

- **Dynamic Profile Selection & Cooperative Transcoding: An Efficient Future for the Multiscreen World** – this paper explores emerging innovations in multiscreen transcoding
- **Improving Adaptive Video Delivery through Active Management** - this paper looks at how service providers can use adaptive delivery technology to deploy unified video processing workflows, which they can use to manage large-scale video delivery over unmanaged networks.