# Cloud-Based DVR and Multiscreen Support Strategies – Optimizing Storage and Transcoding

Ludovic Milin, Carol Ansley
ARRIS

## Abstract

*This paper presents extended and recent usage statistics from in-home multi-room/multi-tuner DVR deployments with an emphasis on characterizing subscriber behavior.*

*These statistics are then extrapolated to formulate sizing assumptions for a typical cloud-based DVR deployment, including ingest, storage and streaming capacities.*

## INTRODUCTION

The advantages of Cloud, or Network, based DVR have been well documented over the past five years. From a cost and operational perspectives, it is perceived as overall cheaper to deploy and operate than in-home DVR, requiring no truck roll to deploy or fix in-home DVR. An additional benefit is the flexibility resulting from a virtually unlimited number of tuners and completely scalable storage space. An nDVR service offer can be changed overnight to add virtual tuners or allow the customer to easily increase the amount of storage they are paying for without any physical changes in the home. Once the storage and playback infrastructure is in place, it also provides easy extension to multi-room, multiscreen services as well as offering potential OTT integration. An operator can seamlessly augment, then replace legacy in-home DVR with new services utilizing the new nDVR resources that can offer a number of ways for the operator to increase its revenue sources, from easy roll out of advanced features like TSTV, Catch-up TV, Pause live TV, to controlling and monetizing advertising on playback, for example preventing ad skipping or replacing existing ads with better targeted alternatives, providing detailed and complete visibility into subscriber playback viewership behavior, and increasing ARPU at high margins for additional recording tuners, disk space, etc.

To date there have been few full scale network DVR deployments and it is difficult for both vendors and operators to accurately predict the amount of storage space or the video ingestion and playback capacities required. While this can be worked around through progressive and controlled rollouts, accurate planning for facilities and operations as well as budgeting for the entire deployment can be difficult. In this paper, we present some study results from a current whole-home DVR deployment with Buckeye that relate well to nDVR considerations.

## MULTI-ROOM/MULTI-TUNER DVR USAGE STATISTICS

The results below come from an analysis of the DVR activity for 1,014 households across 8 representative days in 2013. Each home has 6 tuners available for recordings, and had between 1 and 6 IP set top boxes used to playback the recorded content. Each recording and each playback event was recorded in the utility logs of the system allowing later study and research into user behavior, such as this study.

### Recording Behavior

Two areas can be studied on recording behavior: what did people record, and when did they record it. The first topic is important because of the current state of content licensing regulations. For optimum storage efficiency, a single common copy would be kept of every unique piece of content, but under current regulations and court decisions, a single copy model can only be used if the content provider gives the operator authorization. The choices that people make about what content they are interested in recording affect the relative value of those licensing arrangements to an MSO who is contemplating deployment of nDVR.

The second topic, when do the recordings tend to happen, affects the content ingestion scaling necessary for an nDVR deployment. If the operator must keep separate copies of a piece of content for each subscriber that records it, that affects the overall storage capacity clearly, but it also affects the scaling needed to drive a potentially very large number of recording simultaneously into multiple individual subscriber virtual disks.

Looking first at what content people chose to record, the chart below categorizes the content by unique title. In this view, the dominance of popular content can clearly be seen.

**Table 1**

| | Unique Programs | Percentage of Programs Recorded | Copies | Percentage of Recordings |
|---|---|---|---|---|
| **Recorded once** | 5585 | 40% | 5585 | 9% |
| **Recorded twice** | 2745 | 20% | 5490 | 8% |
| **Recorded 5 or fewer times** | 11889 | 85% | 21804 | 34% |
| **Recorded more than 5 times** | 2128 | 15% | 42896 | 66% |

Slightly more than 14,000 unique programs were recorded across the 1,014 households, but the distribution of those recordings was heavily weighted toward the most popular content. Only 15% of the unique programs accounted for more than 66% of the actual recordings. In chart form, the information can be seen in Figure 1.

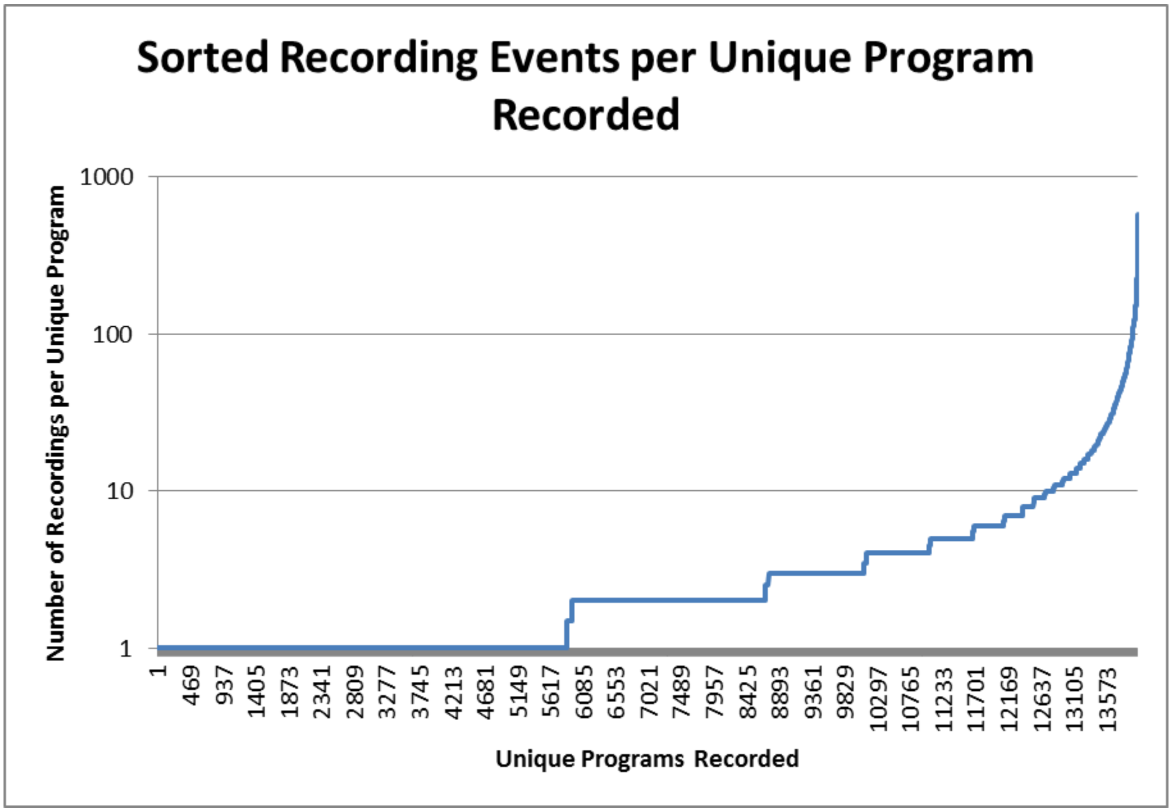


Figure 1

A logarithmic scale is used to make the data more accessible since a linear plot just hugs the axes.

Turning to a consideration of when recording activity happened, the observation that the most popular material dominated the recordings is also reinforced by an analysis of the times when recordings occurred in Figure 2. Recording activity peaks hugely during primetime. Smaller peaks also appear in the afternoon. Since subscribers have the ability to record up to 6 simultaneous programs, Table 2 below will provide details on how many concurrent sessions to expect per subscriber.

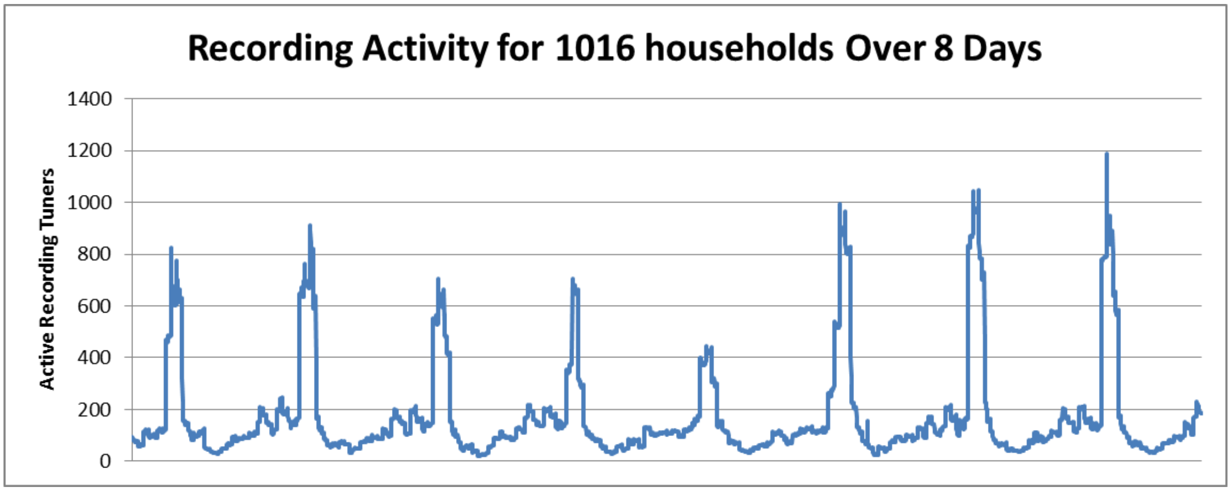


**Figure 2**

If we consider individual household's behavior, we see similar patterns with peaks of usage during primetime, but when individual users are looked at, there is much more seeming randomness to their behavior.

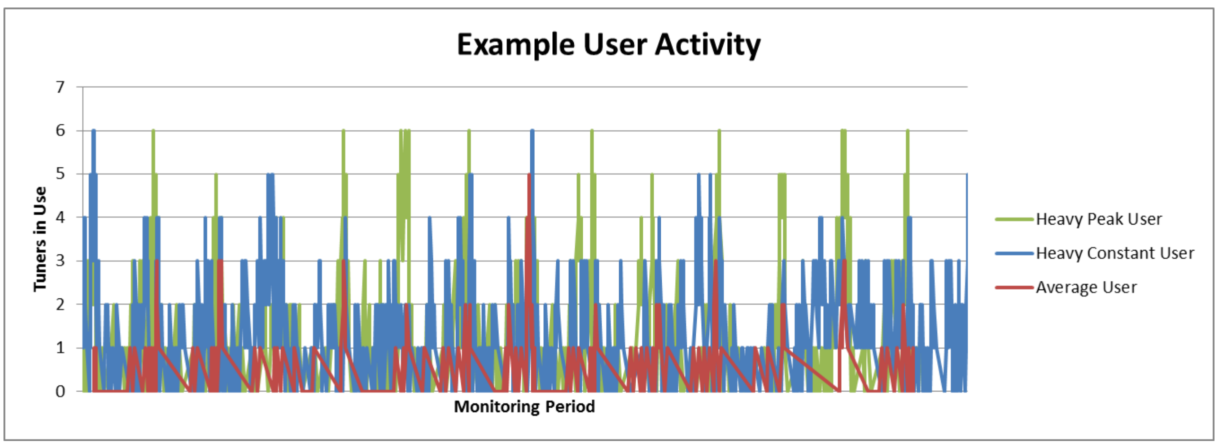


**Figure 3**

These subscribers were chosen as examples of some typical user behavior. The average subscriber tends to record a few programs every day and rarely uses more than 3 tuners. It is interesting to note also that the distribution of subscriber use is skewed with the arithmetic average or mean some distance away from the median value, which is the value in the center of the data. For this data set, we found the following:

**Table 2**

| Number of Sessions | One Tuner | Two Tuners | Three Tuners | Four Tuners | Five Tuners | Six Tuners |
|---|---|---|---|---|---|---|
| **Average** | 49.7 | 24.7 | 11.2 | 4.9 | 1.9 | 0.4 |
| **Median** | 35 | 14 | 3 | 0 | 0 | 0 |

The average subscriber made over 90 recordings over the 8 day period. The spread between average and median implies that the relationship between the heaviest users of the DVR feature and the lightest users is non-linear. When we looked at the amount recorded, the relationship was also seen.

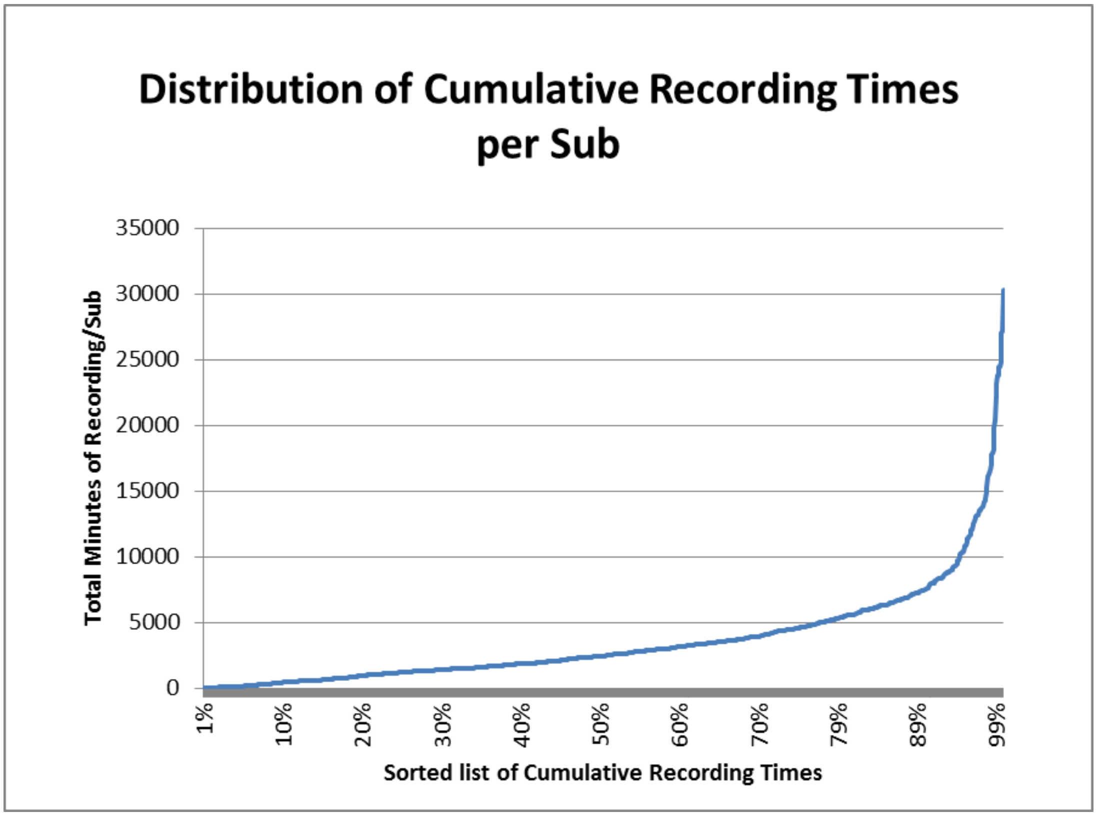


Figure 4

The practical impact of these results is that an nDVR system should be scaled assuming that at least 5 to 10% of the subscribers will use the system to its published limits for active tuners and storage capacity, even though the majority of users will probably use it relatively lightly. On the positive side, about 80% of subscribers recorded fewer than 5000 minutes of content in the 8 day period, which, translated to the worst case of full resolution AVC HD content, only represents approximately a maximum of 300GB of storage. Not that the STB in their possession allows significantly more than this, so there is no bias introduced by hardware limitations. Also note that for those subscribers recording more content over the time period than their STB hard disk capacity allows, they either recorded a significant number of shows in SD resolution, or they deleted content on a regular basis to make room for new recordings.

Playback Behavior

Playback behavior also can have substantial real world implications. The number of streams active simultaneously scales both the nDVR playback infrastructure required to generate and manage the streams as well as the network bandwidth required to carry that video to the end subscribers.

Playback activity also peaked during primetime, with lesser amounts of activity during the rest of the day. Below is a chart showing DVR viewing activity spread across the week.

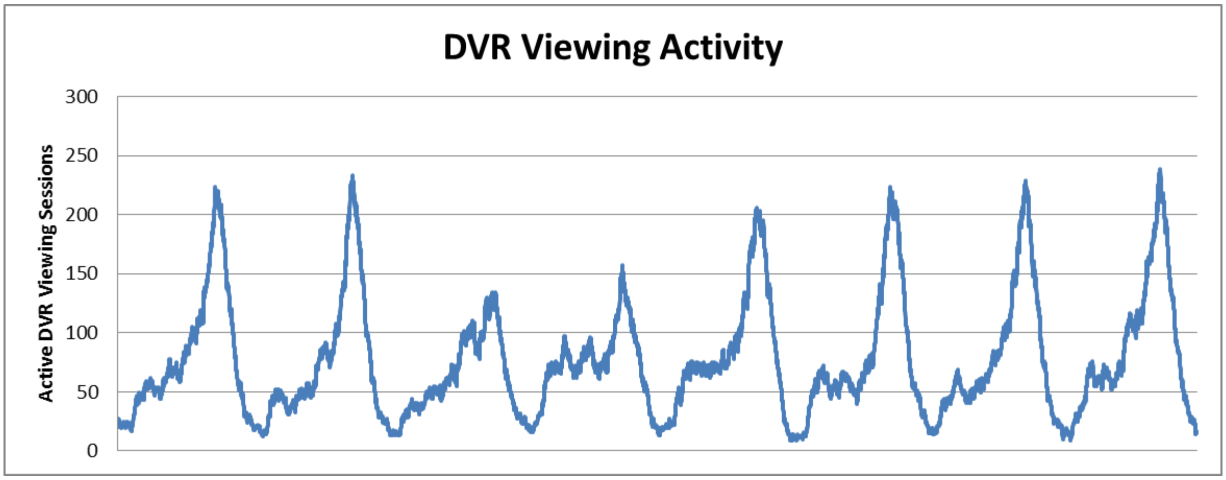


Figure 5

Overall the number of playback sessions varied greatly across the subscribers studied. The majority of the subscribers recorded more content than they played back; overall there were 64,701 recordings made with 53,950 playback sessions for a usage ratio of 77.5%. But, again the individual subscribers' behavior across this week varied widely as shown in Figure 6.

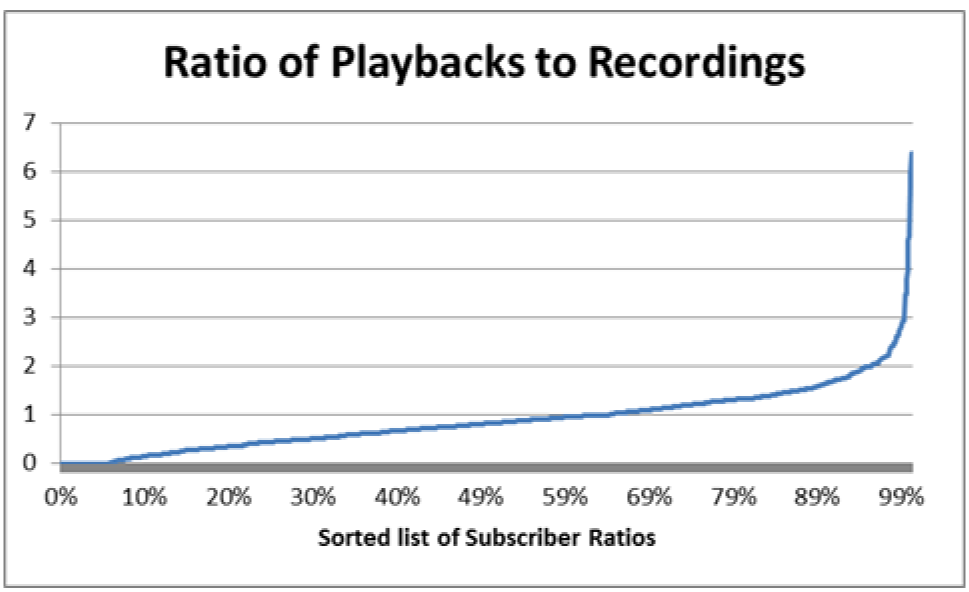


Figure 6

Figure 6 compares the ratios of playback to recordings for the 1,014 households in the study. One issue for this study is that it can take only a snapshot of the activity during the week. Some households showed only recording sessions, others showed only playback sessions – these households were generally light users. Overall about 70% of subscribers record more content than they watch. If the busiest households for DVR activity are considered the ratio of playbacks to recordings is 96.8%.

Age of the Assets

While it is not currently possible to track exactly the time between a recording and its playback or the time between a recording and its deletion, it is possible to report on the number of assets recorded and viewed during the same week. For the period under analysis the number of events in each of the recording and playback categories is shown in Table 3.

**Table 3**

| Events | Total | Unique | Viewed in week | Not Viewed in week | Viewed in week, recorded earlier |
|---|---|---|---|---|---|
| **Recorded** | 64,701 | 14,017 | 41,737 | 22,964 | |
| **Playback** | 53,949 | 11,283 | | | 12,212<br>4,190 Unique |

Fewer than 22.6% of the playback events are for assets recorded over a week prior, and out of those less than a third are for unique content. This would appear to make 7 days a good initial candidate for migrating undeleted recorded assets to archived storage.

Recording and Playback Sessions Duration

The distribution of the recording session duration is shown in Figure 7. As one would expect, two major peaks at 30 and 60 minutes (plus 3-4 minutes of automatic buffer at both ends of the asset) account for almost 70% of all recordings. A minor spike can also be seen at 2 hours.

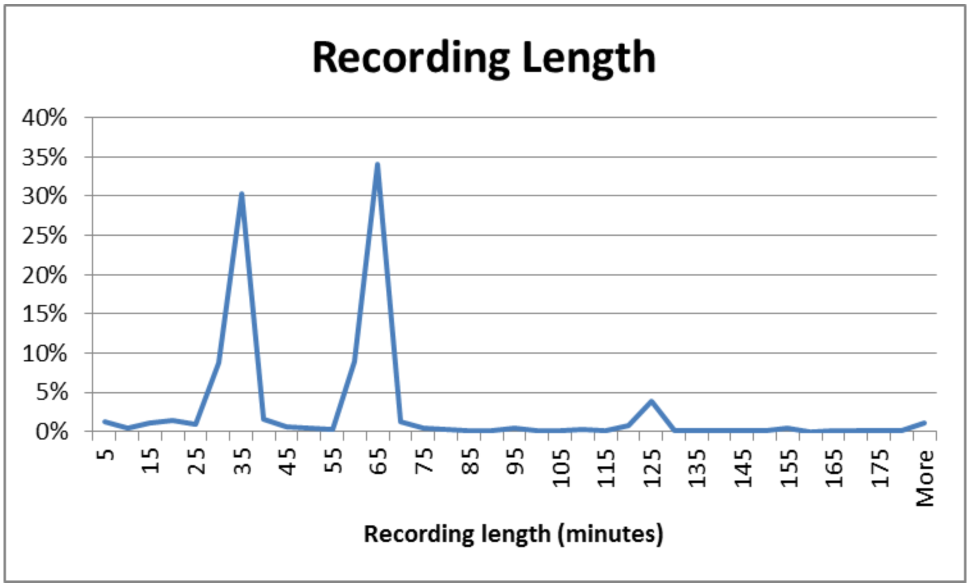


Figure 7

The distribution of playback session duration in Figure 8 on the other end is not so well delineated. It appears that a large number of sessions are abandoned within the first 5 minutes. Then there is a spread to the left of each of the 30 and 60 minutes duration peaks which can likely be attributed to subscribers skipping ads during playback and terminating the session prior to screen credits or similar closing content.

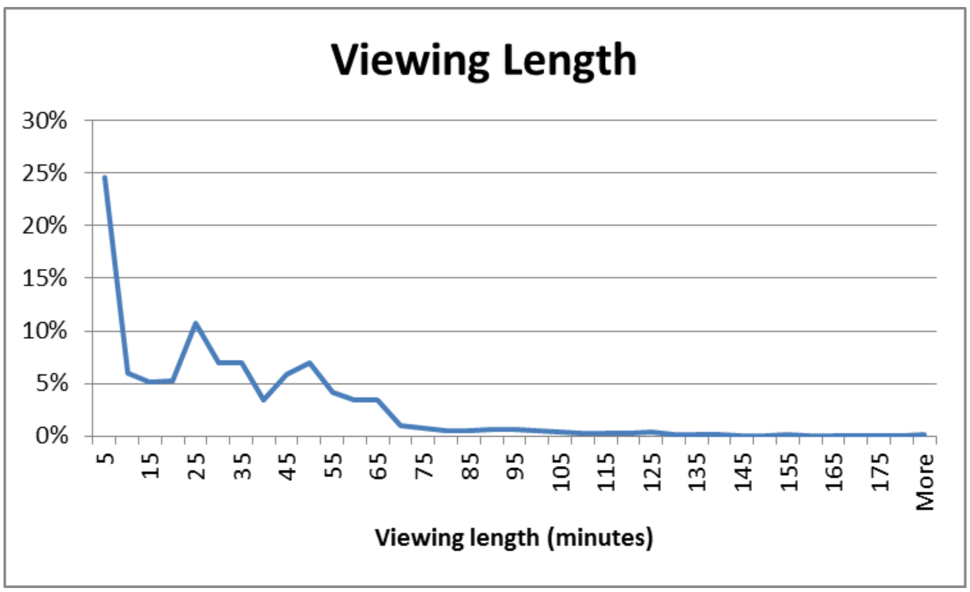


Figure 8

## SIZING A TYPICAL CLOUD BASED DVR DEPLOYMENT

Based on the home DVR data in the previous section, we can propose sizing models for ingest, storage and playback capacity. The initial assumption is that all content is stored as unique/personal copies.

### Storage Requirements

One approach is to size the storage based on the individual quota allotted to each subscriber. However, there are some economies of scale to be had when considering the statistical behavior of a large population of subscribers. As shown in Figure 4, 90% of subscribers typically use less than 400GB of their local 500GB storage. When looking at a large system with a well-tuned archival mechanism, the savings will add up very rapidly and translate either into less upfront CAPEX spending from purchasing 20 to 25% less storage capacity, or additional revenue opportunities by being able to offer increased storage capacity options for a fee to subscribers.

### I/O Capacity

At peak time, Figure 2 shows that there are at least as many active recording sessions as there are subscribers, possibly up to 20% more. From a recording perspective, an ingest system should be designed for being able to record at least 8-9Mbps per subscriber, assuming HD AVC content, and two to four times more when planning multiscreen support with multiple profiles being recorded simultaneously. However, Table 1 shows that most of the content is actually duplicate sessions for popular programs, so the network I/O requirements decrease significantly, by approximately 70% to about 3Mbps, but the burden on content replication and disk writes does not go down.

On the playback side, Figure 5 shows that the number of active sessions only reaches about 25% the number of subscribers; on average this nDVR system should be designed to support an average of about 2Mbps per subscriber disk read and network I/O for HD AVC.

### Content De-Duplication

When moving the recorded content to archive after a few days, seven days appearing to be a good initial setting based on the data in the previous section, one could possibly de-duplicate the recorded sessions and move to a single/shared copy store. There are, however, a few things to remember when planning this: since every subscriber can typically customize the offset for start and stop recording times, every recording of the same show could actually be a slightly different asset, and as a consequence the de-duplication process may not be as efficient as initially hoped.

Based on the numbers in Table 3, with a total of only 14,017 unique assets out of 64,701 recordings, one would expect to save approximately 78% of the required storage space from de-duplication. Even if this is only performed when archiving, the savings should be substantial.

### Adding Multiscreen Support

Multiscreen support typically requires storing and streaming multiple bitrate versions of each asset. More profiles typically provide higher resolution and better video quality where possible, and on the other end also support more degraded network conditions. Additional to the increased transcoding and transport costs, increasing the number of profiles also impacts nDVR storage, ingest I/O and replication capacities, so subscriber experience must be weighed.

We first consider a particular operator five 16:9 AVC profiles use case as described in Table 4.

**Table 4**

| | Resolution | Bitrate (kbps) |
|---|---|---|
| **Profile 1** | 1280 x 720, L4, 29.97FPS | 6,250 |
| **Profile 2** | 1280 x 720, L3.1, 29.97FPS | 3,480 |
| **Profile 3** | 768 x 432, L3.1, 29.97FPS | 1,660 |
| **Profile 4** | 640 x 360, L3.1, 29.97FPS | 1,175 |
| **Profile 5** | 512 x 288, L3.1, 29.97FPS | 940 |

In this scenario, multiscreen support essentially doubles the amount of I/O required from 6.25 Mbps to 13.5 Mbps, and for a 30min asset increases storage from 1.4GB to 3.0GB. A typical scenario with a larger number of AVC profiles is shown in Table 5.

**Table 5**

| | Resolution | Bitrate (kbps) |
|---|---|---|
| **Profile 1** | 1920x1080, High, 29.97FPS | 8,300 |
| **Profile 2** | 1280 x 720, High, 60FPS | 8,300 |
| **Profile 3** | 1280x720, Main, 29.97FPS | 4,600 |
| **Profile 4** | 1280x720, Main , 29.97FPS | 3,000 |
| **Profile 5** | 864x486, High, 29.97FPS | 2,500 |
| **Profile 6** | 864x486, Main, 29.97FPS | 2,000 |
| **Profile 7** | 640x360, Main, 29.97FPS | 1,200 |
| **Profile 8** | 640x360, Main, 29.97FPS | 900 |

In this case, the overhead of multiscreen support is significantly higher, requiring over 3.5 times more storage per asset, 6.9GB versus 1.9GB.

In order to provide the best video quality, but avoid a significant increase in storage requirements, one could decrease the number of profiles offered as the assets get older and moved to archive. Another approach could be to archive only the highest resolution and bitrate version of the asset and use just in time transcoding should there be a request from a multiscreen client. The cost of transcoding is decreasing rapidly with the introduction of Intel I7 GPU based

COTS servers which will soon make just in time transcoding an economical and scalable alternative.

## CONCLUSIONS

Using detailed home based DVR usage data, we were able to provide insight into typical DVR utilization and provide recommendations for initially sizing a cloud based DVR deployment. Most of the content recorded is popular content, limited to 10-20% of the channels, with multiple individual copies, so it is important to consider either negotiating content rights for those channels allowing shared copies, or implementing a de-duplication mechanism when archiving for considerable storage savings. 10% of the subscribers will use 100% of the capabilities given to them but on the other end 80% will use less, some far less, so there is an opportunity for savings by over-subscribing resources. Content older than a week appears to be rarely watched so it can be archived with little impact on subscriber experience.

Continued analysis and staying abreast of existing and upcoming network DVR deployments will help refine these recommendations and allow for local variations of content, multiscreen component, and subscriber behavior. This in turn will enable building stronger and more accurate business cases for nDVR deployments.

## ABBREVIATIONS AND ACRONYMS

| | |
|---|---|
| CDN | Content Delivery Network |
| COTS | Commercial Off-The-Shelf |
| CPE | Customer Premises Equipment |
| DVR | Digital Video Recorder |
| HD | High Definition television |
| HDD | Hard Disk Drive |
| HTTP | HyperText Transfer Protocol |
| MSO | Multiple System Operator |
| RAID | Redundant Array of Independent Disks |
| SD | Standard Definition television |
| STB | Set-top Box |
| TCO | Total Cost of Ownership |
| VOD | Video On Demand |