REMOTE STORAGE DVR

Stephanie Mitchko

Cablevision Systems LLC

Abstract

Approximately seven years ago the team at Cablevision envisioned a system where our subscribers' recorded content could be stored in a data center and played back over our network to their homes. The idea would eliminate the need for costly in home devices and move to a cloud environment. We built a prototype that used commodity hardware along with a combination of vendor and in house developed software that would allow subscribers to create recordings of their favorite shows and store those programs on storage systems in our headend. After announcing our intentions, the content owners filed suit against us, claiming that our system was not permissible under the copyright laws. After several years of litigation, the courts vindicated the system, holding that it did not violate the copyright laws.

In January of 2011 we commercially launched our Remote Storage DVR (RS-DVR) product and marketed it as DVR Plus. Region by region we enabled cloud based services for all of our customers. This paper reviews the overall system and addresses some lessons learned. This will include both technical and operational detail on how moving content recording and playback to the cloud enables new features and portability for our customers. It will also define how this platform is technically extensible to support multiple advanced streaming services as well as Dynamic Ad Insertion (DAI).

Contributors: *Thanks to Peter Caramanica, Rich Neil, Brad Feldman, and John Kenny for their contributions to this paper.*

There are four major areas that are addressed in this paper as part of the overall system,

- Product Features and Flexibility
- Client code architecture
- Software control plane
- Ingest Storage & Streaming

PRODUCT FEATURES AND FLEXIBILITY OF THE CLOUD

Tuners and Storage

The key concept in an RS-DVR solution is moving the storage and tuner functions out of the home. This eliminates the need for storage and tuners in Set-Top boxes and moves the features to the cloud. By removing the need for physical tuners and local storage we break away from the hardware-based restrictions of in home DVRs. Once these features are removed from the physical device in the home, the ability to configure flexible product offerings is technologically available.

For example, feature parameters around storage and simultaneous recordings can be changed and updated easily by utilizing the software control plane. New rate codes control what each subscription tier will look like. A RS-DVR platform can support multiple versions of a DVR product with different storage allocations and tuners. Upgrading storage from 160GB to 500GB or even to 1T can be done by control software only. The same is true for the number of virtual tuners. Since the tuners are not a limitation of the hardware, changing the number is just a variable that is used in the conflict resolution software. Similar to the storage flexibility, simultaneous recordings can be set to four, six, ten, or any number just by utilizing software in the headend.

For example, the initial Cablevision RS-DVR product launched with four tuners and 160GB of storage for the home. As always, any product decision about these values must be carefully vetted to completely understand the impact on the scaling of the complete system. This will be covered in detail later in this paper under the Capacity Planning section.

Whole Home Capability

Moving the recordings to the cloud also makes it possible to provide a whole home solution without adding additional components in the home. When the content is in the cloud, consumers have the potential to access the content from any device in the home. This is not limited to Set-Top boxes as it is now technically possible to deliver this content to any supported device in the home. One could design the system to allow a customer to start watching on her big screen and continue on her iPad.

Any devices in the home could become a DVR by downloading new application software and being authorized for the service. The devices continue to function on the network and do not require any other updates.

CLIENT CODE ARCHITECTURE

In order to enable cloud based recording and playback, client code is needed to utilize the remote systems. There are two areas within the user interface that need to be enhanced to support this architecture, the channel guide and the DVR UI that lists schedules and recordings.

In the channel guide, all record requests are intercepted and directed to the cloud-based RS-DVR backend. For example, if a customer is on the guide and navigates to a program, then selects to record a show from the guide, the information about the program as well as any control parameters that the user might set is sent to the backend software control plane.

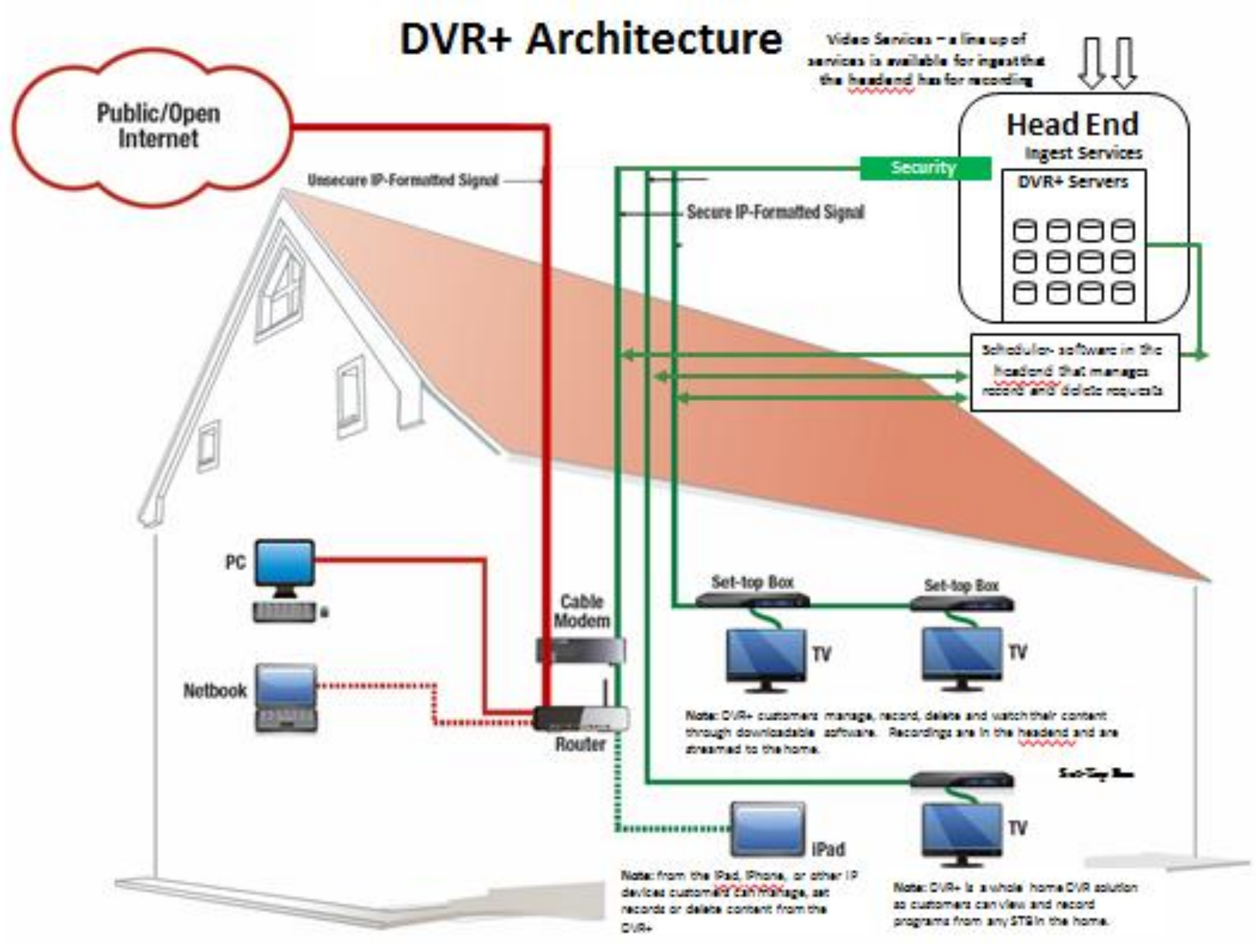

The DVR button on the remote control is programmed to navigate the users' listings of recordings. The RS-DVR guide calls back to the network to retrieve the list of recorded and scheduled programs. Since the DVR guide is in the headend it can be rendered by any STB associated with the account as well as by other IP devices with an Internet connection and proper credentials. Similarly any changes or additions to the requested recordings, new recordings, or deletes are reflected across all devices when rendering the recorded and scheduled lists.

An API layer exists which exposes interfaces to the RS-DVR client application(s). These APIs include functions such as:

- getScheduleList()
- getRecordedList()
- scheduleRecording()
- deleteRecording()

(These are just a few)

These APIs are available and used by all supported devices in the RS-DVR system.

SOFTWARE CONTROL PLANE

Scheduling

The component in the headend that brokers recording requests on behalf of the subscriber is called the Scheduler. The Scheduler is responsible for delivering 'just-in-time' requests to the physical recording system where the subscriber recording occurs.

The Scheduler logically sits between an application on the client device and the systems that control ingest and recording. (Figure 2). Management commands come into the system asynchronously throughout the day. These commands are managed by the Scheduler which, in turn, inform the ingestion service when to start a recording. A well-defined interface between Scheduling and Recording server must exist for this interchange. On successful execution of the recording, the system will commit the recording metadata for display to the subscriber when the RS-DVR menu is invoked.

Series Expansion

This Scheduling component also has an integration point in the headend with guide data. The operation of making recordings is no longer dependent on the STB or even the STB being connected to the network. The Scheduler will use guide data in the headend to make future recordings for series content based on subscribers' requests.

The process of ingesting daily guide data and scheduling recordings based on series settings of the subscriber is a function of the Scheduler. This process is referred to as "series expansion". The function of making recordings is also logically abstracted from the availability of the content RF signals in the home. Details of the ingest services will be discussed in the next section; however, for the purposes of scheduling it is important to point out that content signals are captured at the headend not at the home.

Entitlements

When the recording function moves out of the box and into the headend so does the need to do entitlement checking. Once again the architecture logically as well as physically de-couples the entitlement function from the in-home device. In the in-home DVR, if a subscriber sets a future recording for content they are not authorized for, the STB will tune to the channel and record black (or some barker) since the tuner is not able to decode the signal on that channel. In the cloud all channels are available at the ingestion service, so the system needs to check the subscribed status of the requested content for each subscriber. This software component is part of a standard service delivery platform that is also used for all server-side authorization checks.

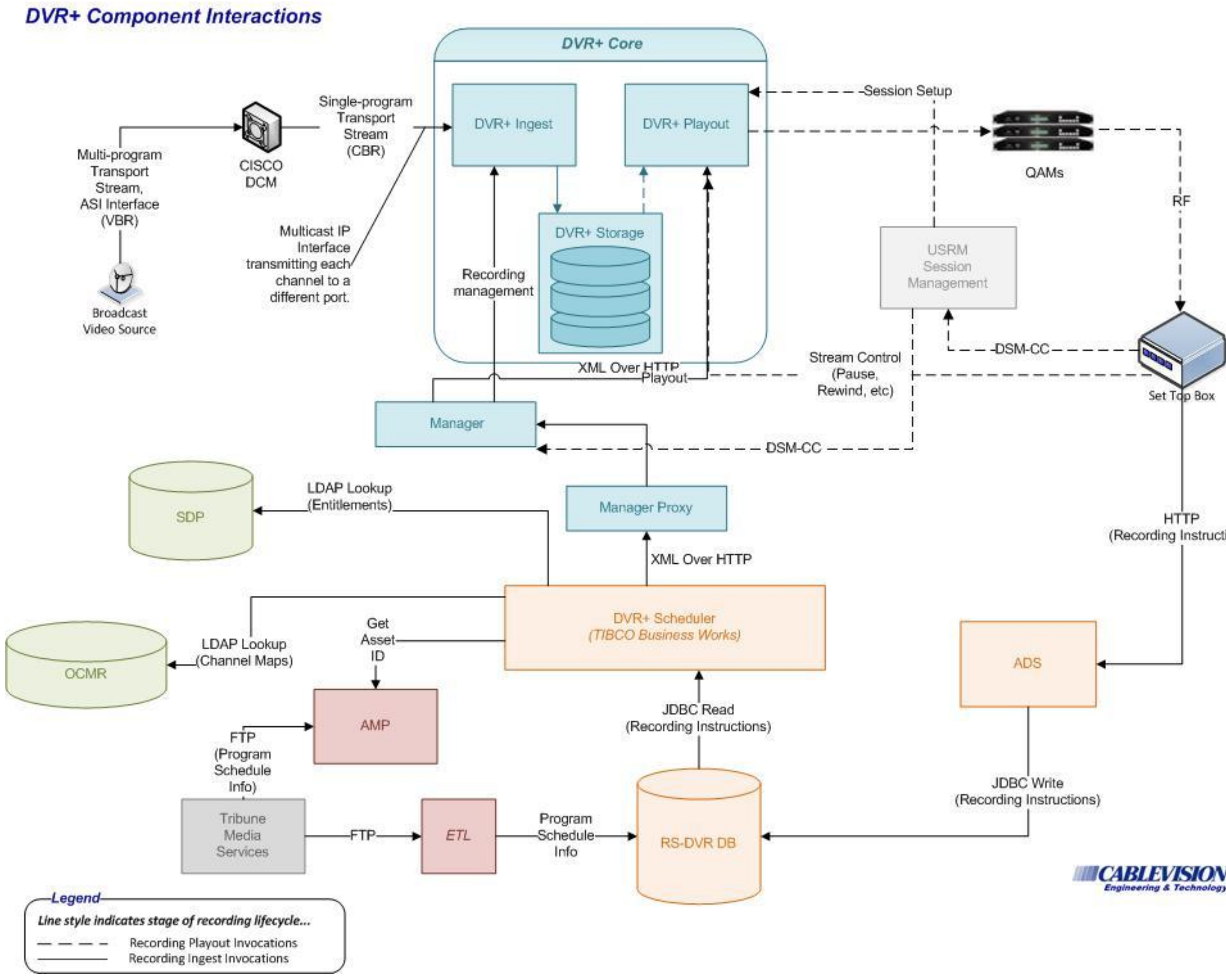

Figure 2 – Reference

INGEST STORAGE AND STREAMING

A main component of RS-DVR is the system that ingests stores and streams the content. (Figure 2. RS-DVR core)

Ingest

Each request received from the Scheduler represents an individual recording command from a subscriber. The ingest process of the RS-DVR system is responsible for ingesting the data/video associated with the program and storing it on behalf of the subscriber.

To enable this, the complete channel lineup must be available to the ingestion process with each channel having a unique multicast IP address assigned to it. The video quality and proper format is critical within this delivery portion of the system.

This component is the network equivalent of the application, in disk based DVR boxes, that tunes the box to the appropriate channel to move the content from the decoder to the local disk. It must be performed in a highly scalable manner as it occurs on behalf of every recording currently scheduled by the system. The data is read from the multicast IP and written to the storage location allocated to the subscriber(s).

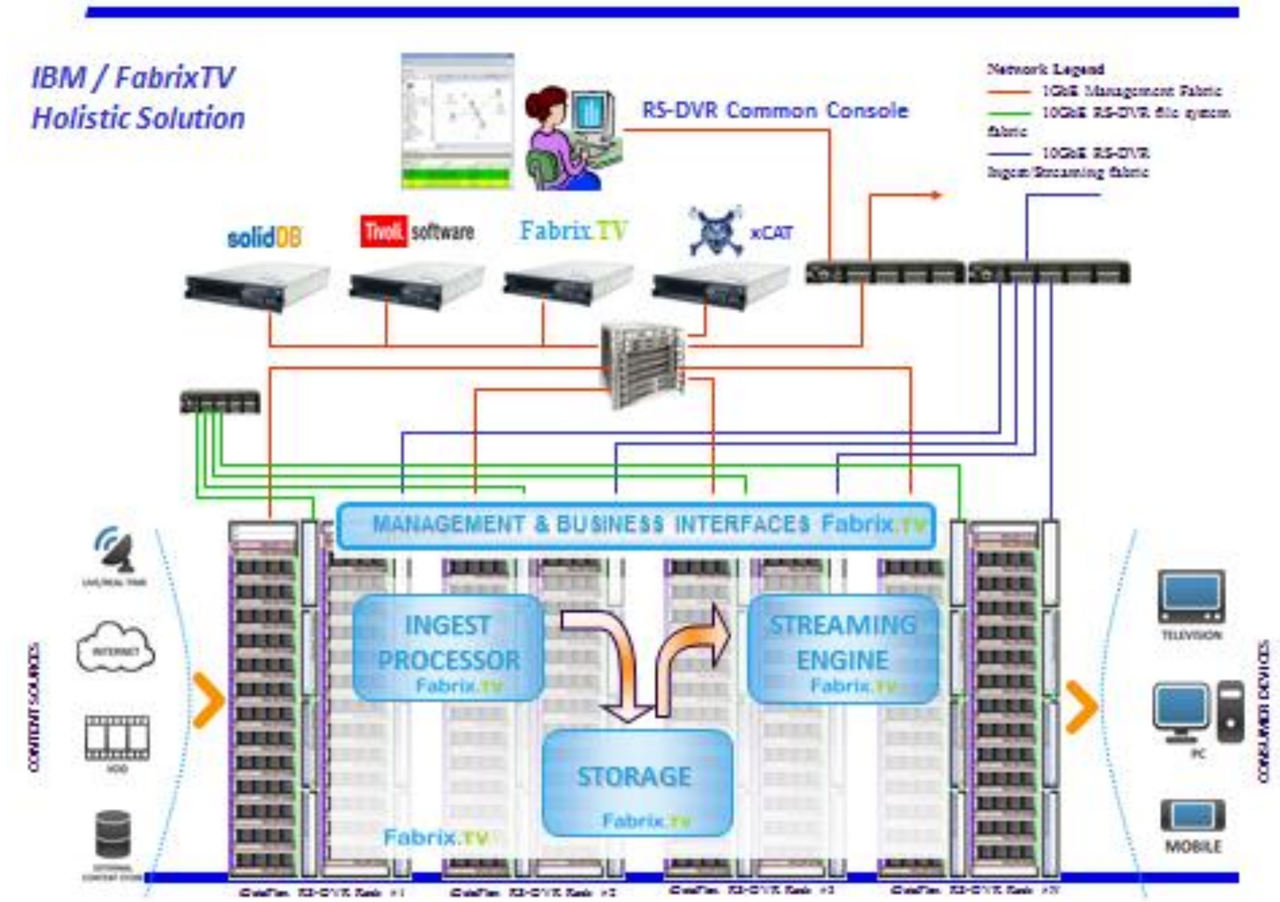

Figure 3 – Operational System Storage

This system demands individual program recordings for each subscriber and requires the right amount of storage that performs well in a write intensive environment. This posed challenges around finding a solution that provides high performance and does not require proprietary hardware. Commodity off the shelf, (COTS) servers and disk storage with the capability to access the disk drives to their maximum level is the solution that works best for this type of system.

Combining the use of a grid-based software file access system provides the best performance to extract the most out of the disk drives.

CAPACITY PLANNING

Building the components described are only a subset of the areas to think about when building the cloud version of DVR. In order to service all requests for recording and playback during peak simultaneous usage of the system it must have enough total throughput capacity in the core. The following is a list of parameters to consider when defining the system scale.

Variables for Storage:

- Total number of subscribers
- Storage allocation per subscriber
- Total available system storage
- Oversubscription percentage

Variables for Ingest:

- Total number of subscribers
- Expected concurrency

System Consideration Parameters

- Disk capacity
- Retro Fit ("Y"/"N")
- Max streaming concurrency
- Max ingest concurrency
- Storage Per Sub Offered
- HD/SD Ratio
- Simultaneous Peak Streams AND Peak Recordings
- HD Rate
- SD Rate

- Homes per Rack (POD)
- Usable Home Capacity
  Product Assumption
- Total Rack Throughput
  Capacity

Assumptions about the usage and product features for the service are used to calculate the system capacity needed to support the customer base. As an example assume the following parameters:

- Subscribers are using 100GB of storage on average
- Concurrency rate at peak usage time 10Mb/s / subscriber

To calculate the number of subscribers that can be put on a reference rack that has 28 servers, 12 drives/server, with drive size at 3T one would use the following:

Calculations of System Requirements DVR+ Core in the Reference Architecture
To go from a single POD to sub count: **Drive Size * No of Drives/server * No of Servers/POD * %age (RAID reserve, cushion, et)** **3TB * 12 * 28 * .85 = 856.8 TB usable storage/POD**
To determine subs/POD based on storage estimates**:** **POD Storage size/Storage Size/Sub (estimated)** **856.6/100GB = 8560 subs/POD**
To calculate total subs to no of PODs based on storage, it's just: **No of Subs*estimated size/POD Storage Size,** **300000 subs * 100GB / 856TB = ~35 PODs**
For ingest capacity**:** **Server ingest capacity * no of servers/POD**
3GB/sec * 28 = 84GB/sec ingest capacity/POD
To determine subs/POD based on peak ingest concurrency: **84GB/sec / 10MB/sec (estimated concurrent) = 8400 subs/POD**
To go from total subs to no of PODs based on ingest**:** **No of Subs*estimated concurrent/POD ingest capacity:** **300000 * 10MB/sec / 84GB/sec = ~36 PODs.**

Other aspects of the system to consider in capacity planning are the following:

<u>Facility Availability</u>

Based on the current demands on storage a facility to house the storage system is required.

<u>Power Requirements</u>

Power in the Headend must be considered as part of the overall installation. In addition, the shift from in-home DVR to DVR Plus in the cloud shifts some of the cost of power from the subscriber to the operator.

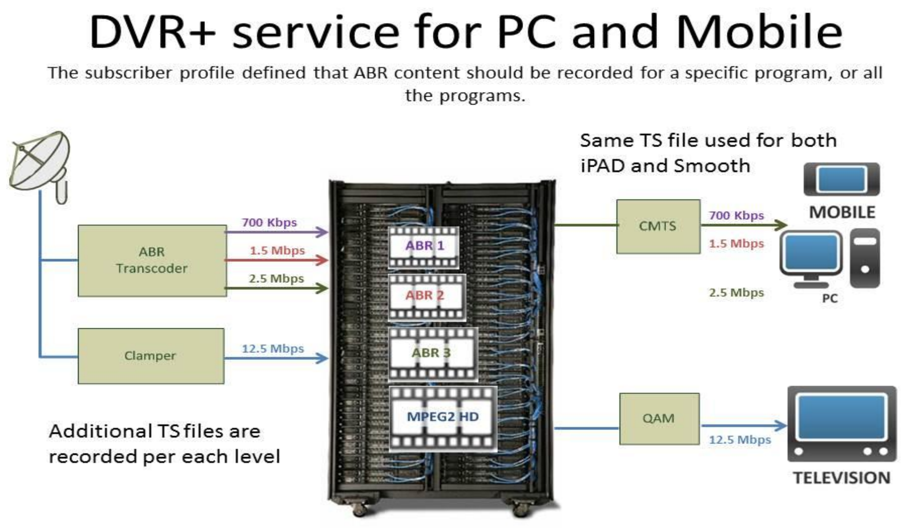

Figure 4 – cross platform distribution

Bandwidth Needs

Considerable network IP bandwidth is needed to deliver data both into and out of the RS-DVR Ingest, Storage and Streaming system. Playback of RS-DVR content to the subscriber uses the same network components as a standard VOD stream. QAM bandwidth is utilized similarly to VOD in that sessions are set up and allocated in order to stream the recorded content to the home. Utilizing a USRM facilitates sharing of the QAM resources between multiple uni-cast applications in this case VOD and RS-DVR. IP bandwidth inside the core network needs to be minimally built to the capacity of the availability of the QAMs. This calculation is based on number of QAMs * the number of service groups * the 38.8 mb/sec bandwidth available at each QAM.

TECHNOLOGY EXTENSIBILIY

The original purpose of moving DVR into the cloud was for recorded content. However, utilizing the RS-DVR cloud system is not limited to just DVR recorded content. There is value in using the cloud for other applications such as VOD, IP content distribution and advanced advertising. The technology component for enabling these types of content is the back office software.

The RS-DVR platform is enabled to support VOD content and library expansion that is not tied to traditional VOD vendor technology. Utilizing the scale of the storage system of the RS-DVR platform it becomes a natural extension of the VOD platform. There is storage available to support expansion of the library and streaming capacity to deliver the content. RS-DVR and VOD technically vary by the mechanism that is used to ingest the content into the system. RS-DVR uses customer-initiated commands to ingest and record content, while the content providers deliver VOD content by traditional pitcher/catcher exchange or over the network via FTP applications that deliver high-speed transfers into the systems. The new technology specific to the RS-DVR system helps in the turnaround time for VOD titles.

The system may also act as an origin server for IP delivered content. The same content sources can be transcoded and packaged within the cloud and available for subscriber requests. Additionally a CDN may be leveraged as the go between for the subscribers and the content.

A dynamic advertising solution is also possible in the platform through the use of SCTE-130 standards. Content ad markers for local insertion opportunities are detected on ingest. On play out, local ads may be inserted in the stream, based on responses from any desired Ad Decision System (ADS). For Cablevision's RS-DVR system, the system maintains linear parity. The ads are the same ads that were shown in the live broadcast within each of the zones. The ad copy is ingested into the RS-DVR core and is available to the streaming system for splicing into time-shifted unicast content as each individual playback calls back to the server for the recording. This advertising solution can be enabled to use any number of ADS systems as potential business emerges.

The original focus of engineering a dynamic advertising solution in the cloud was based on reducing cost and complexity of the overall system. The SCTE-130 solution was engineered to solve the challenge of recording content in multiple ad zones for subscribers in multiple regions. In order to preserve advertising in recorded content without DAI, it is necessary to take multiple versions of broadcast content and make them available to the ingest service of the RS-DVR core. To accomplish this it requires backhauling of the content from regional zones to the master headend. The cost and complexity of this is illustrated by taking the number of local ad zones and multiplying by the number of channels, both HD and SD, used in the commercial insertion system. As an example if there were 100 channels across 36 zones, 3600 video signals would have to be present at the ingest service location.

The DAI solution resolves the complexity and cost of backhauling local ad zones while maintaining parity with the Commercial insertion system. (figure 4.0)

An additional benefit of the cloud solution is the ability to continue to run regardless of the condition of the subscriber's home. For example, if subscribers home is affected due to a disruption services to their home due to a

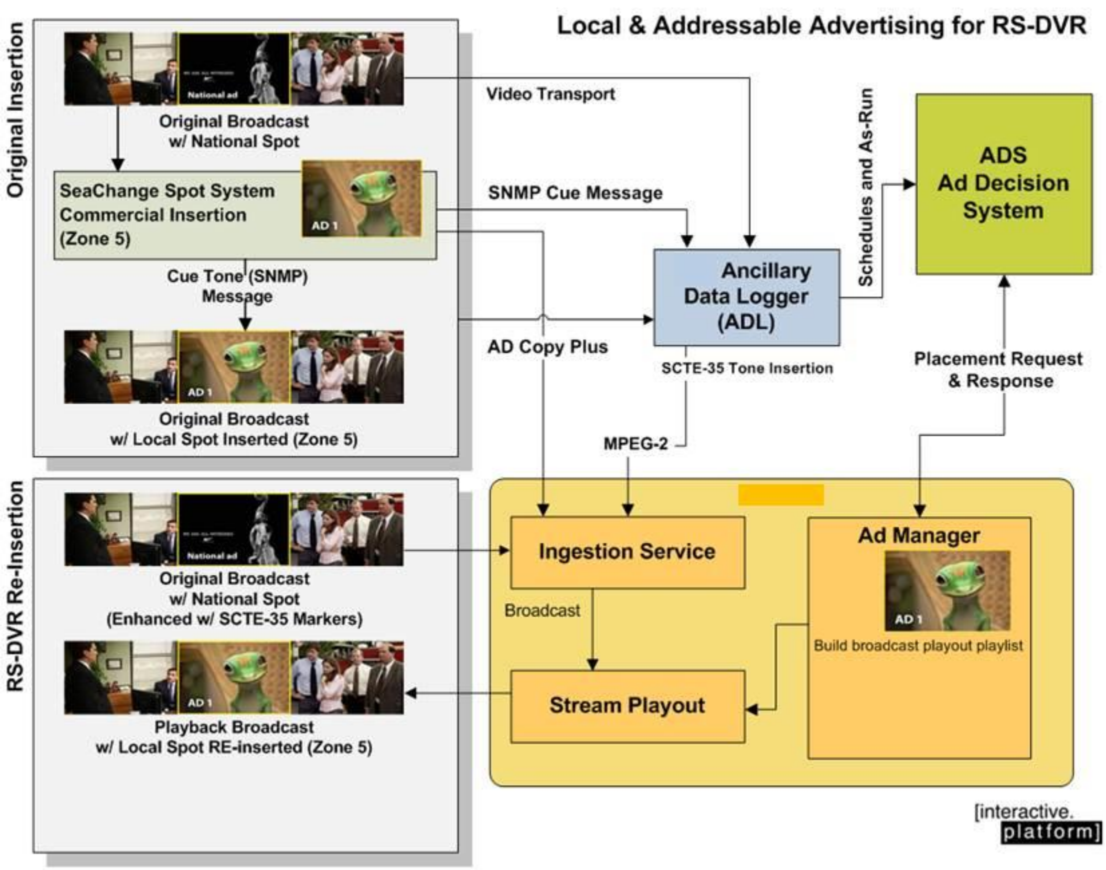

local power outage, their recordings will continue without disruption as long as the DVR Plus system remains online in the headend.

CONCLUSIONS

The idea of moving video services to the cloud is not a new concept. Internet companies are well aware of the fact that having video in the cloud has some strategic advantages. The most obvious one is mobility. Once the content is in a facility that can make it available to multiple devices the product opportunities are endless. Whole home and mobile DVR services are a clear differentiator for the cable companies.

Taking this system from the vision state to deployment is challenging. There are a number of educated guesses that must be made to model streaming bandwidth demands, storage requirements and ingest concurrency. A model is needed to set these values. This model however will only be as accurate as the assumptions that are used in it. After deployment monitoring the system has given us data to validate assumptions. As more subscribers are added to the system it is critical to view and analyze how users are using the product.

Operating this system for two years has given us a unique view into what is needed to support our customers. Since the system operates from the headend there is little for a customer service representative or field rep to do in the home as long as the device in the home, Set-Top box or IP device, is working properly and the home has connectivity. Monitoring for this service is focused on the headend and network components needed to deliver the service. We no longer have to ask the subscriber to return his/her DVR box to a walk in center when the in home disk fails.

Examples of Usage Data

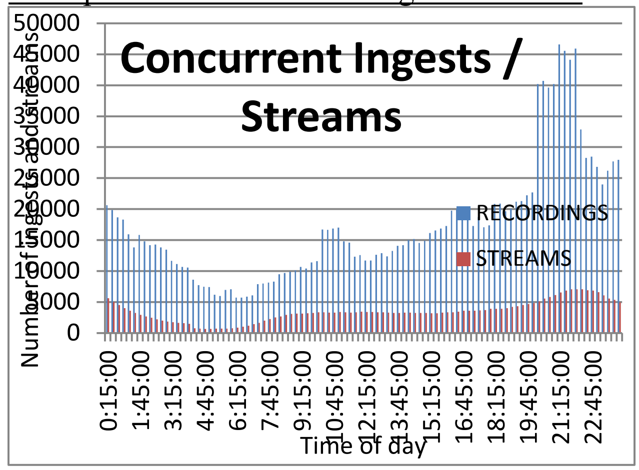

Figure 5 – daily streaming and ingest

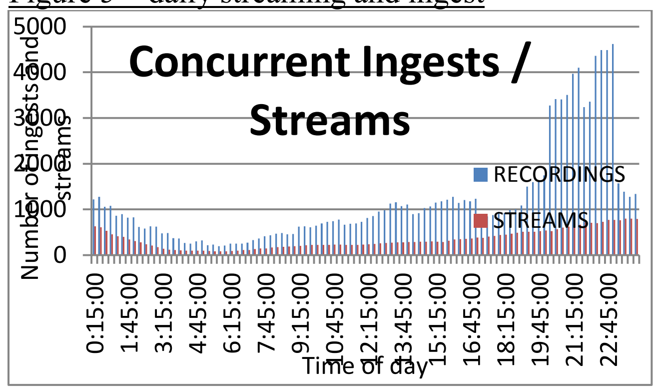

Figure 6 – daily streaming and ingest day 2

The data shown in Figure 5 and figure 6 show streaming and recording activities of a region during a 24 hour period. It shows how many recordings are being made on the system as well as how many streams are initiated. The second chart shows the same data just for a different day. It is clear from the charts that subscribers record more shows than they watch. This pattern is somewhat typical however the usage will vary based on a number of factors including the popularity of a specific show or series.

We must also consider how product features that are offered to the customer in the form of the User Interfaces, preferences, defaults and control affect the performance envelope of the overall system. For example when setting defaults it is helpful to have content delete set to after 14 days with the other options available but left to the subscriber to choose. Conversely we have found that even though we allow customers to record four shows at the same time very few actually do. Figure 7 shows the amount of content that is deleted over time

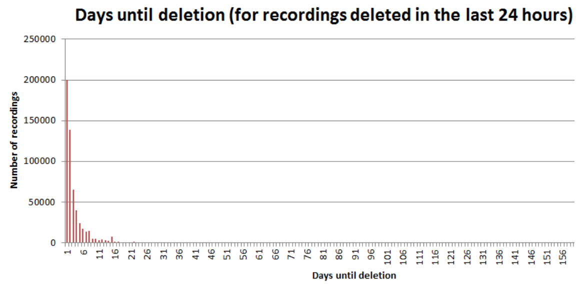

Figure 7 – content deleted over time.

You can see from the graph that the majority of content is deleted less than three days from the time it was recorded with the highest amount deleted in the first day. What this chart also indicates is the many customers keep the default setting for deleting content. It is difficult to see in this chart but at the fourteen day mark there is a noticeable pike in content deletions.

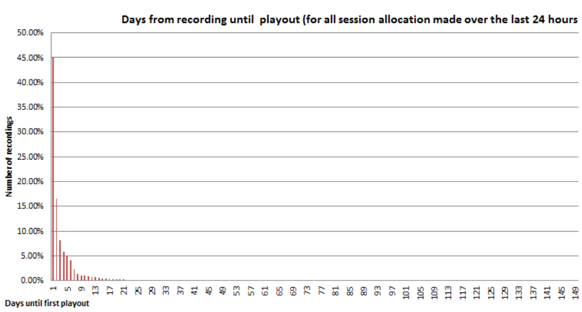

Figure 8 – days until content is watched after recording

This chart shows us that the majority of content is watch within two days from the time it was recorded. So the two graphs above tell us that average behavior for people using the RS-DVR system is that they record, watch and delete content in a relatively short period of time after the show airs.

Figure 9 is a historical view of our RS-DVR deployment. As you can see in the chart streaming and ingest rise at a consistent rate in time as we on board new subscribers to the system.

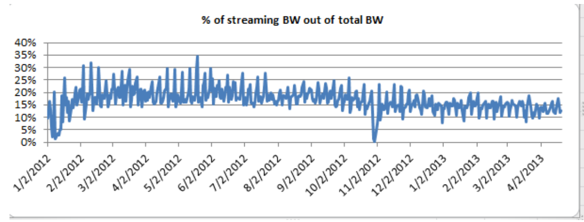

Figure 9 – System Streaming BW utilization

As figure 9 shows it is clear that over time and while the subscriber base is expanding the overall stream bandwidth utilization is low. This is because the system just by the nature of the product is a recording platform. Since the platform has streaming capacity available for other application it is the logical place to use for other content distribution products.

It is important to emphasize, that using a model that takes into account the parameters that most accurately describe the region, number of subscribers and product offering is critical to building out this platform. It is also important to understand the operational ecosystem that a RS-DVR system need to operate in.

This architecture enables the future video grid. It will support the vision of cross platform distribution of content in a manageable way and transform over time from a traditional MPEG2 ingest and storage platform, to a multi-format distribution system that will server up the right format for the client.

Three components of the future architecture are content acquisition infrastructure, a software control plane and a distribution network. Use of the described framework will prepare the industry to continue to deliver quality content and products over multiple networks to numerous devices. Most importantly client/customer owned equipment.