

The Dawn of Cloud-based DVR Services

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Carol Ansley

Sr. Dir. IP Engineering
ARRIS Group, Inc.
3871 Lakefield Dr, Suwanee GA
678-473-8630
carol.ansley@arrisi.com

John Ulm

Fellow of Technical Staff
ARRIS Group, Inc.
22 Topside Rd. Unit 1, Moultonboro, NH
978-609-6028
john.ulm@arrisi.com

Overview

Network based DVR (nDVR) Service has faced a number of legal and economic hurdles over the years. Operator interest in nDVR has accelerated significantly recently. One of the key driving economic factors is the issue of whether a 'unique copy' of content is needed for each and every subscriber recording it or whether a single 'shared copy' can be used for all. In reality, the negotiated rights vary from program to program, so we expect to see a mix of both storage architectures. Since this mix will vary over time, we model these scenarios to show operators the economic impact of offering nDVR service.

The paper continues our analysis and takes a detailed look at the other key drivers for nDVR systems and potential pitfalls. A key concern with capital costs is the trade-off between storage and transport. Operational savings are detailed as learned from early trials. In addition to highlighting these key cost drivers, we take a look at the revenue impact from offering additional services, such as DVR for multi-screen devices; enabling customers to record more than two shows at a time, and extending their content storage in the cloud. These added services map to additional income opportunities for the operator.

To put all of this into perspective, we provide readers with detailed analyses of several key aspects of the intersection between technology and user behavior. We walk you through several scenarios and show economic justification for offering Network based DVR Services. By the conclusion, operators will better understand what it means to offer "DVR in the cloud".

Cloud DVR Nuts and Bolts

The ability to watch “what I want, when I want, where I want” dominates the consumer mindset. Digital Video Recorders (DVRs) introduced the notion of “time-shifting” to consumers years ago. DVRs’ proven popularity has made this a key service for operators to attract and keep subscribers. But this functionality has a cost in capital equipment and operational expenses. As more and more network functionality migrates into the cloud of back office networks, many studies have been done on the best tradeoff between in-home DVRs and DVRs in the cloud. This paper looks at several facets of this topic.

First we compare different instantiations of DVR in the cloud from an economic and functional basis with discussion of the additional service capabilities and revenue opportunities that Cloud DVR offers. The series of legal battles that first brought network DVR to the attention of the industry established a single unique copy per user architecture as permissible. However, as content owners and network operators have gained a more sophisticated understanding of their mutual advantages in offering DVR from the cloud, single shared copy per program based on contractual agreements has become more common.

Next, Cloud DVR network impacts are considered based on some current user behavior studies. Better visibility into user behavior provides more confidence in the predictions of network usage, if Cloud DVR is used as an adjunct to STB-based DVR, or as a replacement for in-home DVR.

Finally, the results are brought together in a case study to show the overall justification for Cloud DVR.

Basic Cloud DVR Components

The basic components of cloud based DVR are shown in Figure 1 below. A video data plane processes and delivers the content, driven by a control plane that manages the video content’s work flow.

Walking a piece of content through the process begins with the Ingest and Grooming functions. Video transcoding and transrating is usually required to ensure that all of the required resolutions and codecs are available to the Recorder for storage. The Recorder component represents all of the network storage which may be centralized or distributed. The Streamer component enables the real time delivery of video content upon a user’s request.

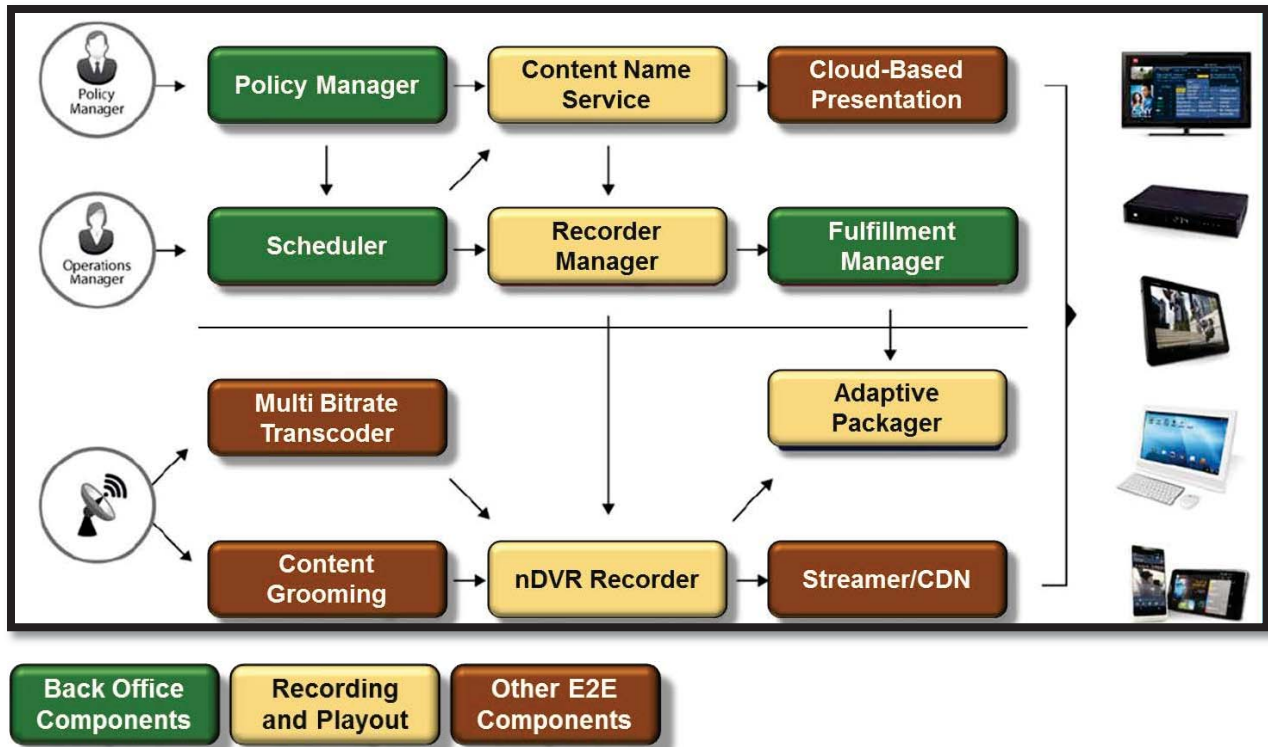


Figure 1 – Basic Cloud Based DVR Components

The control plane allows simple or complex management of the video content. The Scheduler component collates requests from users to produce a recording schedule, and may also incorporate predictive algorithms to select the programs to be stored as well as the resolutions and codecs to be used for each recording. The Policy Manager component allows the operator to set recording policies including which programs may or may not be recorded, how long recordings can be held for a user, and potentially which recordings are targeted to the network storage facilities and which recordings may be held on a local DVR drive, if available. The nDVR Content Management component takes inputs from the Policy Manager and Scheduler to manage the overall content storage and retention. The Recorder Manager component uses information from the Scheduler and the Content Manager to direct the Recorder component. The Fulfillment Manager interacts with the Streamer and the Packager component, if needed, to supply content to the end users when they request it.

The architecture pictured above has a lot of flexibility and supports a variety of recording and streaming models. It can store and deliver Cloud DVR content in multiple formats to an IPTV environment, a legacy QAM system, an IP video system via a Content Delivery Network (CDN), or to a hybrid system with both QAM and IP devices. A key advantage of this approach is that it provides the flexibility to operate as a centralized library for on-demand content, as well as supporting centralized ingest with edge streaming. Network considerations dictate which solution is optimal, but in an nDVR system with a mix of Time Shifted TV (TSTV) and unique copy content, edge cache is almost always

indicated. Figure 1 provides a generic illustration that includes both a centralized on-demand content store and as an edge streamer in a hybrid QAM/IP system.

The above architecture is based on the trends seen at operators today addressing the multi-screen nDVR challenge. A centralized Recorder captures all content in one centralized library, potentially including the multiple bit rate content prior to packaging. This approach leverages the efficiency and scalability of the Recorder by centralizing the recording and storage content destined for both operator-owned STBs and multi-screen uses. The advantages of this approach include:

- Simplifying management by allowing centralized storage of all content including VOD, nDVR content for STBs and nDVR content for multi-screen devices.
- Reducing storage requirements by consolidating to one copy of each SD, HD and multiple bit-rate file (per recording location) when content agreements allow it. Various adaptive bit-rate versions can be packaged when requested, meaning only one copy of each bit-rate version is stored to support multiple adaptive streaming formats (Apple HLS, Microsoft SmoothStreaming, etc.)
- Leveraging high-density recording platforms to reduce hardware, power, cooling and operational costs.

Another option is an edge transcoder and packager. For systems with a relatively small subscriber count but large content library, this model may be a viable option.

Recording/Capture

As directed by the Cloud DVR Scheduler, the Recorder captures HD, SD (MPEG-2 or H.264) and/or multiple bit rate H.264 transport streams along with the appropriate metadata to support subscriber access of stored content across both STBs and mobile devices. Note – depending on the content management solution selected, the Scheduler could be a single component or separate components.

When a Recorder is deployed in a centralized manner to support multiple streaming sites, the number of copies of a particular VOD asset that must be captured and stored is substantially reduced, greatly lowering storage costs for fixed content. The centralized recording approach does not require, but offers the option of deploying edge streamers that can cache VOD or shared copy nDVR content at the edge of the network to reduce network traffic. By using a mezzanine format such as MPEG-2 Transport Streams rather than packaging every copy in several different formats (Apple HLS, Microsoft SmoothStreaming, Adobe, etc.) the amount of storage can be greatly reduced. By using MPEG-2 TS as the mezzanine format, theoretically the same bitrate can be used in a multi-screen environment for STBs as well as second and third screen devices like tablets, though in practice this may not be realistic.

The Scheduler sends requests to the Recorder Manager, which ideally manages a scalable cluster of Recorders, allowing the Unique Copy Recorder infrastructure to scale to multiple petabytes (10^{15}) of storage for hundreds of thousands of subscribers in a manageable way.

The Recorder then serves as an origin server for requests for both STB-based content (from the QAM and IP edge streamers) and for multi-screen content (using a Just-in-time packager as the origin interface). The Recorder may also serve as the VOD library for QAM and IP Edge Streamers if an operator chooses to centralize VOD storage as well.

STB Delivery

In a centralized model, the Back Office sets up sessions on the appropriate Cloud DVR Recorder cluster, which can stream and provide trick modes to STBs. Alternately, as shared copy Cloud DVR libraries grow, it may become more efficient to cache at least the most popular content at the edge using Streamers. The back office would set up sessions at the Streamer based on the requestor's service group. The Edge Clusters retrieve content from the Recorder as needed. Ideally, the Recorder clusters should support RTSP-based streaming as well as HTTP-based byte range transfers to allow the use of edge Streamers where permitted under content use agreements, and to allow direct support for unique copy play-out where necessary.

An intelligent Streamer might not require trick files, so that only the content and index file need to be delivered to the streamer, reducing cache size and network bandwidth usage. Importantly, the Streamer efficiently requests byte ranges of only the needed content and generates trick files on the fly at the edge of the network. This eliminates the need to send an entire file to the edge Streamer over HTTP to support trick play.

Multi-Screen Delivery

Alternately, the IP Streamer can retrieve content via a CDN. The CDN can view the Recorder as an origin server, or the Cloud DVR Recorder can publish assets to the CDN Origin. In this model, the Transport Stream files destined for Set-top Boxes (STB) are also cached in the CDN, along with the Adaptive Streaming fragments for three screen devices. However, this results in duplicate copies of STB-based content in the ecosystem.

The Packager can generate the appropriate Adaptive Streaming format in real-time based upon the request of the client. These fragments are then cached within the CDN, eliminating the need to retrieve them again from the Packager and Recorder. By packaging on the fly, only one central storage format (H.264 in an MPEG-2 Transport Stream) needs be captured and stored in the Recorder for each bit rate. Upon packaging, multiple Digital Rights Management formats are supported, depending on the type of client and Adaptive Streaming protocol.

Economics of nDVR

DVR architectures can generically be split into in-home and cloud-based DVR. In-home DVR architectures consist of STB-based or Gateway-based DVR. Cloud based architectures include unique-copy Remote Storage DVR (RS-DVR), and shared-copy network DVR (nDVR). Each has its own advantages and disadvantages.

In-Home DVR Architectures

STB-based DVR

STB-based DVR was the first model to be deployed and has been around for years. It consists of a STB with local disk storage. As Hard Disk Drive (HDD) prices per bit have dropped, this model has served the consumer and operator fairly well. It has offered an easy to deploy model with ever expanding storage capabilities as the STBs have followed the PC's HDD integration curves. Operators are quite comfortable with this operational model.

From a consumer's perspective, the size of their storage is fixed once the STB is deployed. Also, the STB design often limits consumers to only be able to tune to two programs at any instance in time. So, the consumer can record a maximum of two programs simultaneously, but only record one program if watching another. Consumers are also reluctant to trade in their DVR for fear of losing their saved programs.

From an operator's perspective, STB-based DVR has several disadvantages. The presence of a HDD in the STB increases its failure rate, and when a HDD fails, the STB must be replaced, sometimes with a truck roll to the consumer's residence. Many STBs did not offer the operator the ability to track DVR viewing at first, which reduced the number of ad views for programming that was recorded. From a revenue perspective, while most operators were able to charge additional fees for DVR STBs, the advertising embedded in the recorded programming was hit twice in value. Most DVRs offer the option to skip advertisements, and even if the ads are viewed, they are often no longer relevant at the time the subscriber chooses to view the recorded material.

Gateway-based DVR

The Home Gateway is a newer generation of consumer premises equipment (CPE) devices that operators are now deploying. The Home Gateway includes the HDD and acts as DVR for the entire home. It communicates with other CPE devices such as additional STBs distributed through a home network such as MoCA or WiFi.

The GW-DVR brings several advantages over STB-based DVR. There is only one HDD per home, although it may need to have a larger storage capacity. The Home Gateway also has more tuners, with six tuners being a typical number. While a couple of these tuners might be used for live viewing, there are more available for recording than with STB-DVR.

The GW-DVR does have many of the same disadvantages as the STB-DVR with respect to limited disk storage, failure rates and lost advertising revenue.

Cloud based DVR Architectures

Remote-storage DVR

RS-DVR offers the operator the option to store recordings in a centralized storage facility in the headend, or even farther back in the network. This model was approved

by the U.S. courts as a simple replication of the STB-based storage and requires a separate unique copy of content for each and every subscriber. RS-DVR offers some advantages over the STB-based model, but has some challenges of its own. Within the headend, the operator can share HDD capacity among users, though not sharing individual programs. This facility offers the end user better reliability since the cloud storage can be set up in a redundant array to prevent data loss due to any single HDD failure. To the operator, RS-DVR also offers reduced capital costs for STBs, though with an increase in headend power and space usage. The removal of the HDD from the STB also increases its reliability, since the HDD contributes significantly to unit failures over its lifetime.

Having the storage in the cloud also provides a new revenue opportunity for the service provider. The provider can now charge the consumers more for additional storage space and the ability to record more than two programs simultaneously. It also enables other services like “Whole Home DVR” to all devices inside the home, even if there is no home network. Finally, they can deploy a DVR service to subscribers with older STB without having a truck roll to install the DVR inside the home.

A challenge to the operator, though, is the increased bandwidth needed to supply the recorded content to the end user. For STB-based DVR, the programming is for the most part already broadcast to the end user, so bandwidth consumption patterns are not changed by its deployment. With RS-DVR, new unicast bandwidth is needed to supply the subscriber with their recorded content when they are ready to view it.

Network Impacts of RS-DVR

To explore the impacts of remote-storage DVR on the network, the DVR usage habits of about 1000 DVR subscribers were studied across one week. Each DVR household had 2.5 STBs on average. The following graph shows DVR recordings and DVR views using one minute accuracy for that week.




Figure 2 – DVR Recording and Viewing Behavior

The DVR viewership results show that, on average, more than one fifth of all households were viewing DVR content during prime time. This statistic indicates that the narrowcast bandwidth dedicated to a service group must be sized to allow more than one fifth of the households to be able to access DVR narrowcast content simultaneously. It is worth noting that this result does not include real-time pause capability, also called Personal DVR services. Users who paused and resumed their live content would generate traffic in excess of that shown here.

Using an example Serving Group with 500 homes passed and 250 subscribers, a minimum peak allowance would need to provide for at least 50 active nDVR sessions. With a 50% mix of HD and SD traffic being delivered as H.264/MPEG-4, the operator might need almost 200Mbps, or 5 - 6 additional narrowcast channels, to support RS-DVR. Operator's mileage may vary quite a bit depending on Serving Group size, encoding bit rates and HD/SD mixes.

Network DVR

An optimized DVR solution would only store a single shared copy of each program instead of a unique copy for each user. This solution requires the content provider's consent, since it does not comply with the original remote storage DVR court ruling. In the remainder of this paper, nDVR will be used for common shared copy per program solutions, and RS-DVR for single unique copy per user solutions. The differences in network storage requirements resulting from only having to store each piece of content once are substantial. To usefully compare the economic impact of the two cloud based solutions this paper uses the same example service group with about 1000 DVR customers. In the analysis of recordings, it was found that if a program was recorded, it was recorded about five times. This means that the RS-DVR solution would require approximately five times the storage of the shared copy nDVR solution.

Mixed RS-DVR and nDVR Solutions

In reality, operators will have a mix of content with the shared copy rights to some programs and other programs requiring a unique copy per user. And the ratio between these will be changing over time. What is the impact of this mixed solution?

First, the statistic above that a recorded program in the 1000 user study was recorded five times on average is not the whole story. That statistic obscures the larger power-law relationship that was seen. The most popular 10% of programs that were recorded constituted almost 60% of all recordings and were recorded an average of 29 times each. The most popular program was recorded 579 times. The remaining 90% of programs were recorded about twice each, with 42% of the unique programs recorded only once.




Figure 3 – Basic Cloud Based DVR Components

As an operator contemplates approaching content providers for single shared copy storage rights, this consumption data shows that the rights for single shared copy storage of popular content are quite valuable in terms of reduction in storage costs, while the right to store single shared copies of less popular content is much less valuable. The data highlighted above demonstrates that having the shared copy rights to the most popular 10% of programs would save the operator more than 60% of their cloud storage requirements.

Economic Analysis of DVR Models

To examine the potential economic impact of different DVR models, several different factors should be considered. The capital outlay and total cost of ownership of in-home DVR and cloud-based DVR have different factors but can be estimated. The in-home DVR costs factor in the added costs of the HDD, plus an estimate for additional home networking costs to distribute content within the home.

Cloud based DVR has several infrastructure costs to consider. One is the impact on the HFC network of the additional unicast traffic used to bring network DVR video to the end users. Other additional costs of cloud-based DVR include additional edge QAM ports, additional transport costs, as well as the costs of the equipment dedicated to DVR support in the network (the Ingest equipment, the Recorder, the management infrastructure, etc.) For nDVR as well as RS-DVR, the initial capital costs outside of storage infrastructure are similar. A substantial advantage does accrue to nDVR-based

architectures for storage costs, however. The findings of the study described in the previous section show that nDVR can offer a substantial advantage. The following graph shows an example comparison of relative capital costs on a per-subscriber basis. This does not include operational costs such as truck rolls, which are discussed later.




Figure 4 – Relative Capital Costs for DVR Deployments

Another relationship to be considered is the advantages brought by scale for nDVR compared to RS-DVR. As the number of subscribers served by a centralized nDVR system increases, the amount of storage increases at a much slower rate than for RS-DVR. With RS-DVR, storage needs scale linearly. This is shown in the figure below.




Figure 5 – Basic Cloud Based DVR Components

Total Cost of Ownership

The total cost of ownership (TCO) for DVR systems is more complicated to express in a standardized way than the initial capital costs because of the variety of choices available to an operator. Perhaps the biggest contributor to operational costs is the need to do a truck roll to the customer's premises. This may occur at either initial installation or if problems occur later. Cloud based DVR solutions have a potential advantage in that they can eliminate truck rolls for initial installation for customers who already own a STB without a HDD. Cloud based DVR solutions also allow operators to use simpler CPE devices that could enable more consumer self-installations.

The TCO must also consider the increased failure rate of a DVR STB because of the HDD or other storage media. A STB failure may also result in a truck roll for some operators, while others may encourage the subscriber to return the unit to an MSO location for a replacement unit. The failure rate of enterprise class HDDs used in headend storage locations is also usually lower than the drives used in DVR STBs. Additionally, most headend systems will use some form of redundancy such as RAID (Redundant Array of Independent Disks), so that any individual drive failure will not result in lost content and can thus be replaced on a regular maintenance schedule, versus on a rush or emergency basis, easing the operational concerns of network storage.

While the capital outlay for GW-DVR and nDVR as shown in Figure 4 above are comparable, nDVR will have a clear advantage in TCO.

Additional Cloud-based Revenue Streams

A unique advantage of the cloud-based DVR system is that advertisement avails in the original content may be updated with fresh material suited to the time of the playback and the subscriber. For example, a limited-time sale ad may be replaced with a more current one. Inserts of upcoming "News at 10" might be refreshed with the latest relevant news clips and highlight the next upcoming locally-produced new show. Advertisers may be willing to pay more money for targeted ad avails that are more likely to cause the viewers to watch them. The operator could also disable or reduce the ad-skipping functionality of the DVR content, but without a concurrent ad replacement feature, that choice could generate subscriber dissatisfaction.

Other revenue sources may come from "up selling" more storage and features such as unlimited simultaneously recording. For some operators, nDVR may be a way of offering a complete video service suite to all multi-screen devices, such as tablets, smartphones, gaming devices and smart televisions. For a complete economic analysis, operators need to look at their service offerings and estimated revenue impacts of nDVR as well.

The Pragmatic Choice for the Future - Flexibility

Few operators are currently in the position to be able to completely convert to DVR in the Cloud. A significant deployed base of DVR STBs exists in most operators'

networks. Some of these units may be able to work in a cooperative way with a Cloud-based DVR solution, but at a minimum, a significant firmware update may be needed. Additionally, few operators have the large amounts of unused spectrum necessary for a quick rollout. A practical path for the deployment of a Cloud DVR system may be a hybrid approach that allows older DVR STBs to access the resources of the Cloud, but not switch to it fully immediately.

Also, even though not all content owners may be willing to allow common shared copy nDVR systems, many may be willing to discuss that option. A Hybrid system that can accommodate both shared copy and unique copy content with flexible business rules is a necessity to allow a gradual rollout of Cloud-based DVR systems that can fit into a practical rollout envelope for network bandwidth and capital concerns.

Conclusion

The dawn of Cloud-based DVR systems is already upon the industry. The basic features are well established, but now the focus is on finding the best deployment models for different operators with different demographics and subscriber bases. Economic justification for Cloud-based DVR comes from several directions: capital costs, operational costs and considerations, as well as revenue opportunities. The short-term simplicity of STB-based DVRs runs aground on the longer term considerations of increased STB failure rate and its lackluster reputation among advertisers.

Cloud-based DVR, whether shared copy nDVR, unique copy RS-DVR, or a mix, offers improved long-term economic models and the possibility of increased advertising revenue. Full nDVR implementation may be some time in the future as some operators and content providers are still cautious over the court rulings, plus some content providers' reluctance to adopt shared copy models. This paper recommends that operators select a flexible Cloud DVR solution that can offer customized business rules per content provider. With the dawn of cloud-based DVR, the future looks bright for increasing customer satisfaction while also putting DVR on a better economic footing.

Abbreviations and Acronyms

CDN	Content Delivery Network
CPE	Customer Premises Equipment
DVR	Digital Video Recorder
HD	High Definition television
HDD	Hard Disk Drive
HFC	Hybrid Fiber Coax
HTTP	HyperText Transfer Protocol
MSO	Multiple System Operator
QAM	Quadrature Amplitude Modulation
RAID	Redundant Array of Independent Disks
SD	Standard Definition television
STB	Settop Box
TCO	Total Cost of Ownership
VOD	Video On Demand