# The Future of IP Video: From Pay TV to Cloud TV

Dear reader,

When Scottish inventor John Logie Baird demonstrated the very first televised moving images in 1926, he could scarcely have imagined how users would consume TV today. But even taking the terrestrial, cable and satellite TV eras into account, more dramatic changes have occurred in the last ten years of television than in the previous 80 combined. And all this change was made possible by one thing: IP video.

Today, consumers watch video on a variety of connected devices and demand a superb experience on all of them. New over-the-top (OTT) providers such as Netflix® are offering direct-to-consumer services with low prices, advanced user interfaces and easy access to multiscreen video. Changing usage patterns brought on by subscriber desire to watch content at the time, location and on the device of their choosing are increasing content distribution costs. Pay TV providers are particularly susceptible to these trends and need to adapt their traditional TV delivery architectures to offer innovative services that attract and retain customers.

So what will television platforms look like in five or ten years? At Alcatel-Lucent, we predict the following trends:

The traditional set-top box (STB) will disappear. The functions of today's STB hardware will be carried out in the network and by the connected device itself, eliminating the cost and complexity of managing home-based STBs.

Traffic will be all unicast. Over time, device format fragmentation, time-shifting viewing habits and service personalization will erode broadcast and multicast efficiencies. Ultimately, every end user will be served with a unique stream.

Services will be deployed in the cloud. Dedicated video platforms will migrate to cloud-based services, reducing costs and accelerating time to market. Operators will move from vertically integrated middleware stacks to more open architectures with best-of-breed components.

In the following articles, we'll dive into Cloud DVR, DNLA CVP2 and Content Delivery Networks: three key technologies that pay TV providers need to embrace to migrate smoothly to IP video and realize the potential of cloud TV.

We hope you find these articles informative and welcome your feedback. For more about Alcatel-Lucent's vision, strategy and solutions for IP video, please visit **www.alcatel-lucent.com/ip-video**.

Best wishes,

**Paul Larbey**

President of Alcatel-Lucent Video Business Unit

# Clear Skies Ahead for Cloud DVR

By Roland Mestric


**HIGHLIGHTS**

- Cloud DVR technology makes all TV content available on demand, on any device and in any location
- Service providers reduce the cost of delivering DVR services while increasing features for subscribers
- New revenue opportunities are enabled across the TV content creation and delivery ecosystem

Moving video storage and processing from the set-top box to the network benefits subscribers, service providers and content owners alike.

## SOMETHING FOR EVERYONE

Moving digital video recorders (DVRs) into the cloud is the natural next step for home-based DVRs. Service providers can use cloud-based DVRs to reach more subscribers, devices and locations with content. Cloud DVR technology removes hardware – such as hard disks and tuners – from set-top boxes (STBs) and puts it into the service providers' network (Figure 1). This move improves end-user experience and reduces cost and complexity in the service provider's network. It is part of a broader **IP Video strategy**, in which service providers can use the network to **deliver premium content** directly to a broad range of consumer-owned devices.

With **cloud-based DVRs**, end users can watch their favorite TV shows on any device, anytime, anywhere. Service providers can deliver DVR services without installing or maintaining expensive hardware in subscribers' homes, which reduces costs. Furthermore, service providers can increase profitability by offering new capabilities to both subscribers and partners.

In addition to opening new revenue opportunities for service providers, cloud DVR technology extends benefits to other members of the TV ecosystem (such as advertisers, broadcasters, TV channels and content owners), who gain greater control over how content is recorded and consumed.

Despite these benefits, content rights legislation has not kept pace with technology in all jurisdictions. Local regulations may affect how service providers can offer cloud DVR services to their subscribers.

## IMPROVED VIEWING EXPERIENCE FOR SUBSCRIBERS

### TV on demand on any device

The cloud DVR approach allows end users to access the TV programs they have recorded from any device, with full DVR capabilities. They can also access a range of time-shifted TV services, including pause live TV, rewind TV, start-over TV and catch-up TV. Cloud DVR technology allows subscribers to start watching on one device in the home and continue on another while on the go.

Subscribers can also use cloud-based DVRs to record programs while away from home, even if the STB is turned off.

### Unlimited storage

Moving to the cloud means DVR storage capacity isn't limited by the size of the STB hard disk. Subscribers can potentially access unlimited hours of video from the cloud.

### Parallel recordings

Recording in the cloud does not consume any access bandwidth or any STB processing power, allowing subscribers to simultaneously record multiple TV programs while watching linear TV.

**Figure 1. The cloud DVR principle**

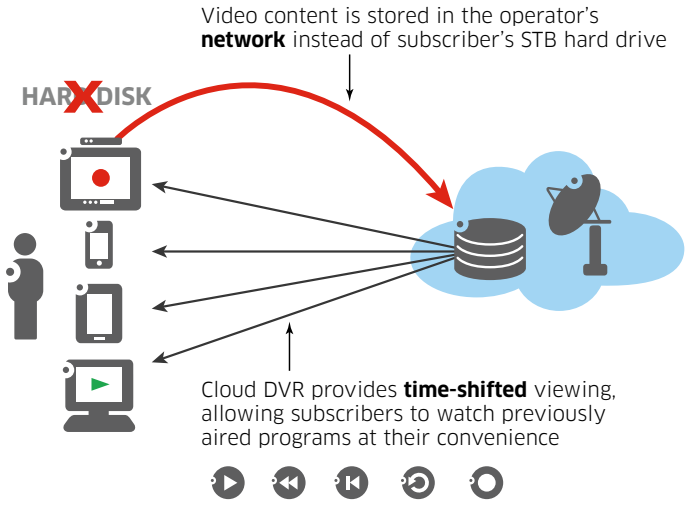

### Reliability and availability

Content stored on an STB can be lost if the hard disk crashes or if the STB has to be changed. Service providers can build their cloud DVR services with redundant storage that guarantees content is always available. The solution can also provide some back-up capabilities for subscriber personal content.

## REDUCED COSTS AND INCREASED REVENUES FOR SERVICE PROVIDERS

### Lower capital expenditure (CAPEX)

DVR components – including the hard drive, multiple tuners and additional memory – make up a large portion of the cost of DVR-capable STBs. Moving these components to the cloud greatly simplifies STBs. Using central storage is less expensive than the end-to-end cost of maintaining large numbers of hard disks in home-based DVRs, particularly when the "shared copy" approach can be used (see the content rights challenge section).

Service providers can deliver DVR services to any STB in subscribers' homes, and don't need to deploy any other hardware to support whole-home DVR capabilities. A subscriber can turn any "smart" device into a DVR by downloading an application and authorizing it for the service.

### Reduced operational expenditure (OPEX)

Cloud-based storage makes it easier to operate and maintain DVR services, particularly when hard disks fail. The increased reliability and availability of storage in the cloud reduces the number of calls to a service provider's help desk, as well as hardware shipments to subscribers' homes, with or without truck rolls.

### Incremental revenue

Service providers can potentially increase revenues by offering their subscribers value-added, time-shifted TV services that are available on all devices.

Cloud DVR technology allows service providers to charge for increased storage capacity and for the ability to record more programs at the same time. They can make these extensions happen in the network without making any changes to end-user equipment or shipping new hardware.

Cloud DVR deployments may not create ew revenue opportunities in every market. Nevertheless, the technology will improve subscriber experience and the perceived value of the service provider's TV service, increasing customer satisfaction and reducing churn. Cloud DVR technology can also kick off a multiscreen or over-the-top strategy, which will attract both new and existing customers with engaging premium services.

## BENEFITS FOR THE ECOSYSTEM

### Ads are still watched with DVRs

Subscribers and service providers are not the only ones to gain from cloud DVR technology; advertisers, broadcasters, TV channels and other content owners can also benefit.

Advertisers initially expressed concern about DVRs when the technology came to market because they feared subscribers would skip commercials. Studies show, however, that viewers watch ads both at the time of broadcast and on programs they record with their DVRs, effectively increasing viewership[1]. This shows that the business model for broadcasters and advertisers is not entirely broken by DVRs.

Time-shifted viewing methods have largely complemented traditional TV business models and subscriptions, rather than acting as a substitute. Another study from the UK shows that DVR-equipped homes watch more hours of TV in total than those without DVRs, and this effect can compensate for any ad-skipping that does occur[2].

Finally, research has shown viewer's attention actually increases while fast-forwarding commercials[3] and in response, advertisers have created ads that work equally well if subscribers watch them live or fast-forward them with their DVR[4].

### Increased control

Unlike home-based DVRs, cloud DVR technology allows participants in the TV value chain to control variables such as storage capacity, how long recordings can be stored and the number of tuners available to record content at the same time. The technology can also specify the type and number of mobile devices subscribers can use to playback content.

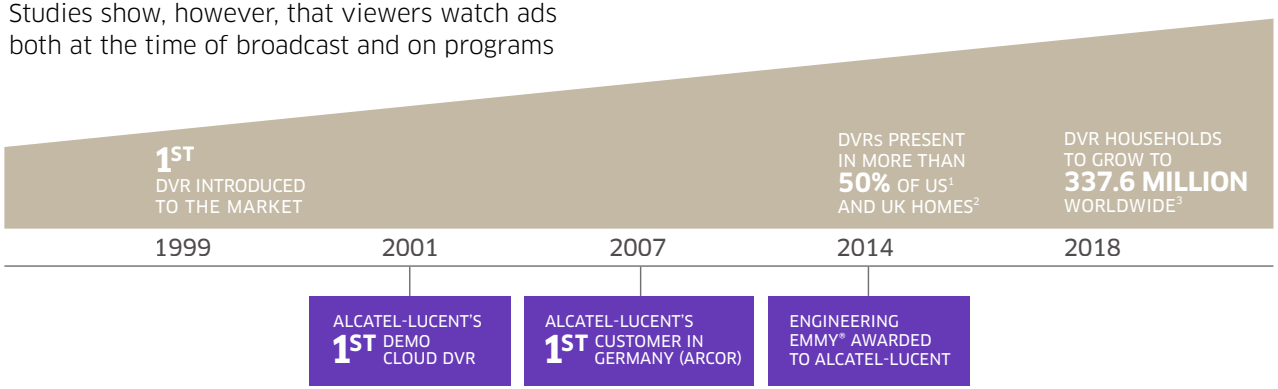


1. Nielsen: **"No hardware, no problem: VOD lets users time-shift with ease"**, September 2013.

2. Streaming Media Europe: **"Thinkbox shows growth in U.K. viewing on TVs and mobile devices"**, February 2013.

3. Global Industry Analysts: **"Global Digital Video Recorder (DVR) Households to Reach 337.6 Million by 2018, According to New Report by Global Industry Analysts, Inc."**, July 2012.

### New revenue opportunities

Cloud-based DVRs also give service and content providers more insight into how consumers watch live and time-shifted TV, potentially opening the door to new revenue models. For example, they can use this information to better target their content and services. An advertisement might not be appropriate if a subscriber watches a recorded program weeks or months after it first aired. Advertisers can replace commercials in catch-up and recorded programs with new ones targeting individuals or groups of users, making advertisers more relevant and providing additional revenues for broadcasters and service providers.

## THE CONTENT RIGHTS CHALLENGE

Despite the benefits for all players in the value chain, content rights issues have inhibited cloud DVR deployment in many regions.

While copying video content or TV programs on personal devices such as DVRs is generally accepted, it's a different story for recording and storing TV programs on the service provider's network. Whether or not cloud-based recording violates copyright laws has significant implications for service providers.

The model service providers use to deploy their cloud-based DVR services (Figure 2) depends on agreements negotiated with content providers, as well as country-specific intellectual property rights legislation.

### Shared-copy cloud DVR

Shared-copy cloud DVR deployments store a single copy of each recorded program in the network and use them for multiple subscribers. The key advantage is efficiency: the service only ingests and stores one copy of each recorded program, significantly reducing storage costs.

### Private-copy cloud DVR

Private-copy cloud DVR deployments store one copy of the requested TV program in the network per subscriber, and each stored program is only available to the subscriber who made the recording request. If two end users request the same content to be recorded, duplicate copies will be made, making this model less cost-efficient than the shared-copy model.

## MATCHING TECHNOLOGY TO CONTENT RIGHTS – TODAY AND TOMORROW

Cloud DVR technology needs to be flexible to allow service providers to fully benefit from its implementation. Ideally, service providers will be able to use the shared-copy model, as they have done in Switzerland. In less favorable situations, such as in the United States, service providers can deploy the private-copy model. This approach will ease the technology's acceptance by content owners or comply with copyright legislation, and provide a path to a more efficient solution. Service providers can then introduce shared copies on a-per content basis using simple software settings as negotiations with content owners progress. This would allow service providers and subscribers to **benefit from the flexibility and efficiency of cloud DVR technology** today and in the future.

## FOOTNOTES

1. **Nielsen, "C3 TV Ratings Show Impact of DVR Ad Viewing"**, April 4, 2009.
2. **Thinkbox, Op. cit.**
3. S. Adam Brasel & James Gips, "Breaking Through Fast-Forwarding: Brand Information and Visual Attention", American Marketing Association Journal of Marketing, 2008.
4. **Daniela Walker, PSFK, "VW Ads Target Viewers Who Fast Forward Through Commercials"**

**Figure 2. Cloud and home-based DVR models for storing content**

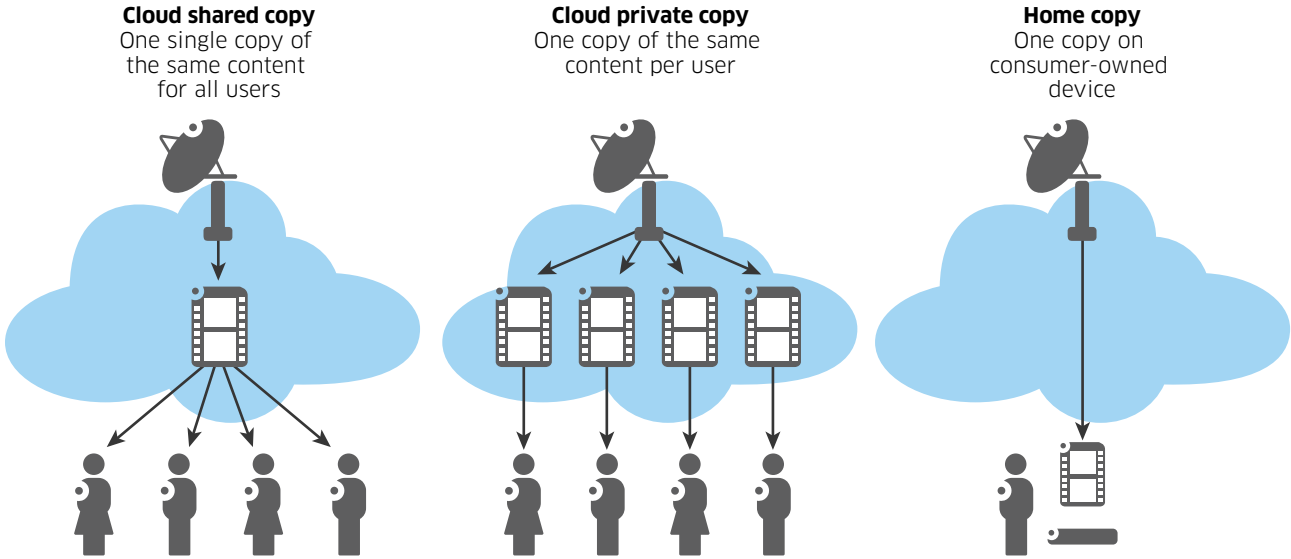

# The Cloud Delivers Premium TV on More Devices, Faster

By David Robinson

By Christy Thomas

**HIGHLIGHTS**

- DLNA CVP2 enables pay TV operators to securely stream premium content to multiple devices
- Operators can quickly bring video streaming services to IP devices by combining technologies like HTML5 remote user interfaces and MPEG-DASH
- Operators can remove cost and complexity from the home by using CVP2 to migrate capabilities to the network and stream content from the cloud

Consumers want to watch premium content on every connected device in the home. To maintain their supremacy in the home, pay TV operators need the ability to quickly make their video services available on every new device that reaches the market. By capitalizing on Digital Living Network Alliance (DLNA) standards, including the new Commercial Video Profile 2 (CVP2) specification, operators can gain a fast and cost-effective means to reach more devices with premium content from the cloud.

## NEW SCREENS BRING NEW OPPORTUNITIES

Today's fast broadband networks make it possible to stream video to consumers on demand in the home. This increased speed has enabled new online video streaming services – Apple, Netflix, Amazon/LoveFilm, Hulu and others – to capture a share of the home entertainment market. It has also inspired broadcasters to extend their offers to devices beyond the TV, as seen with TV Everywhere in the US and YouView and SkyGo in the UK.

Not to be outdone, pay TV operators are exploring **IP Video** as the next way to deliver TV services to the home. With IP video, operators can use the network to deliver premium content directly to a broad range of consumer-purchased devices. Network-based delivery will remove the need for in-home digital video recorders (DVRs) set-top boxes (STBs) and complex home gateways, which currently cost operators billions of dollars each year.

## NEW SCREENS ALSO BRING FRAGMENTATION

Device fragmentation presents a formidable obstacle to aspiring multiscreen pay TV operators. These operators need to meet the demand for premium content across a fast-growing and increasingly diverse collection of devices, including smart TVs, tablets, smartphones, PCs and game consoles. To provide a captivating multiscreen experience, pay TV operators need to deliver on expectations in three key areas:

- **Content discovery and control:** Deliver interfaces that make it easy for users to find, select, record and consume content across a variety of screens. Smartphone and tablet applications can provide these functions, but they need the ability to recognize and interact with other devices in the home.
- **Content streaming:** Ensure that all in-home devices can support content streamed as unicast IP traffic. Adapt formats and bit rates to different devices and available bandwidth.
- **Content protection:** Authenticate subscribers using IP addresses, username/password credentials or device information. Scramble streamed content using per-session encryption keys or a full DRM scheme. Restrict sharing to the home network.

Two factors raise the degree of difficulty for operators. One is the lack of similarity between smart TV operating systems. The other is the concurrent use of many variants of the Android operating system. To deliver a superior IP video experience to every subscriber on every screen, an operator must support hundreds of combinations of devices and operating systems.

## DLNA OFFERS SOLUTIONS TO THE MULTISCREEN CHALLENGE

The DLNA CVP2 specification offers solutions to the **device fragmentation challenge**. It provides a quick and cost-effective means for pay TV operators to extend whole-home DVR and multi-room services across many devices. Moreover, it can empower operators to use their networks to stream premium content to subscriber-owned devices on demand.

DLNA specifications focus on ensuring interoperability between certified wired and wireless devices. They eliminate the need for operators to develop clients for specific functions like device interoperation and control. The CVP2 specification adds capabilities like auto-discovery, HTML5 support, client security, content encoding and adaptive streaming. Pay TV operators can use these capabilities to quickly develop universal clients that reduce the complexity and cost involved in extending their services throughout the connected home.

### Device interaction and content discovery

DLNA allows devices on the home network to seamlessly find each other, discover each other's capabilities and interact. For example, a TV or game console can get and stream content by connecting to a network-attached storage device or PC acting as a media server. Or, a smartphone or tablet can act as a remote control that helps users discover content on a media server and play it on a TV.

### Enhanced user interfaces and HTML5 support

Connected consumers expect consistent and compelling user interfaces (UIs). DLNA is addressing these expectations by building **HTML5 support** into the CVP2 specification. The addition of HTML5 will enable operators to create remote UIs that take advantage of the web browsing capabilities and superior processing power offered by the current generation of connected devices.

By using web-based HTML5 UIs, operators can extend a consistent user experience across a broad range of devices. There is no need to develop a separate application for each device/OS combination. HTML5 web pages download the user experience to the client device. If required, operators can make quick updates to the look and feel of these pages and make them available to all devices at once.

### Content streaming through MPEG-DASH

Most CE devices purchased today include built-in support for **HTTP Adaptive Streaming** (HAS) technologies that can adapt the streaming bit rate to the available bandwidth.

With HAS, content is encoded into multiple bit rates. Clients dynamically pull the bit rate that is best suited to current network conditions. DLNA supports Dynamic Adaptive Streaming over HTTP (MPEG-DASH), a standardized form of HAS. It also supports other forms of HAS, including Microsoft Smooth Streaming and Apple Live Streaming.

### Cloud-based content sourcing

Pay TV operators' current DLNA implementations rely on home gateways that act as media servers. A gateway receives content from the operator's network, processes it as needed and delivers it to other devices in the home. This approach is a logical choice for MSOs that run television services over QAM networks and require QAM tuning by way of a device in the home.

Home gateway-based media servers make little sense for telcos that run pure IP networks, or for MSOs that are migrating to all-IP networks. These operators can reduce costs by migrating key media server capabilities and processing functions – content preparation, transcoding, recording, UI rendering and device-specific packaging – to the network.

Migration may be a gradual process for cable operators, with clients rendering unified UIs from media server sources both in the home (high usage linear QAM-based video channels) and from the network. Over time, sources will migrate from the home to the network and reduce reliance on home gateways and STBs.

**Figure 1. Streaming premium content to connected devices in the home with DLNA CVP2**

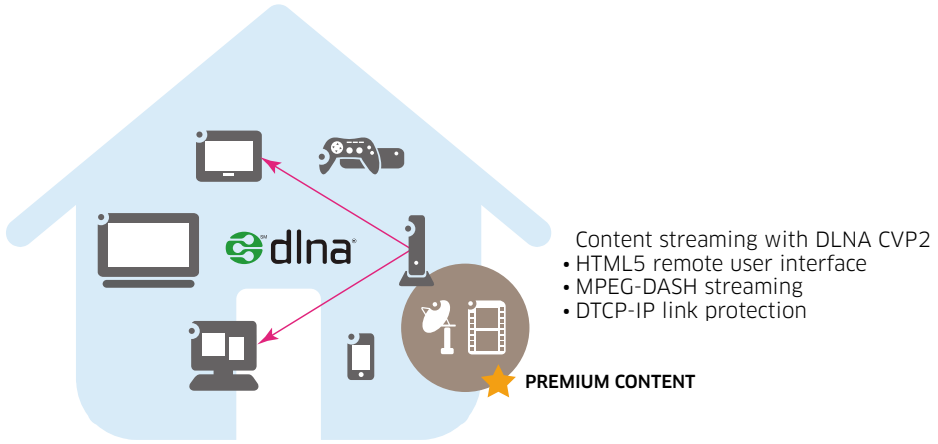

### Content protection

DLNA specifies that Digital Content Protection over IP (DTCP-IP) be used for moving content between devices. DTCP-IP allows any pair of devices to establish a secure link with no input from a digital rights management (DRM) system. To restrict content sharing to the home environment, DTCP-IP imposes localization rules that govern time to live, round trip time and concurrent source streams.

These rules exist because DTCP-IP security is weaker than most DRM schemes. It is adequate for in-home sharing of media between DLNA-compliant devices. But an MSO or IPTV operator knows that it is delivering to a subscriber on its network. It is therefore logical for the operator to treat DLNA sources as extensions of the home network and work with TV studios and the Digital Transmission Licensing Administrator (DTLA) to relax these localization rules. Relaxed rules would allow the operator to stream directly from its network to devices in subscribers' homes.

## SHAPING THE FUTURE OF PAY TV

The DLNA CVP2 specification will enable cable and IPTV operators to quickly and cost-effectively extend cloud-based streaming TV services to multiple devices in subscribers' homes. By allowing operators to migrate processing functions to the network, it will help create a simpler connected home that is free of set-top boxes. Operators that take advantage of CVP2 will deliver pay TV services directly to subscribers and avoid having to ship hardware to their homes.

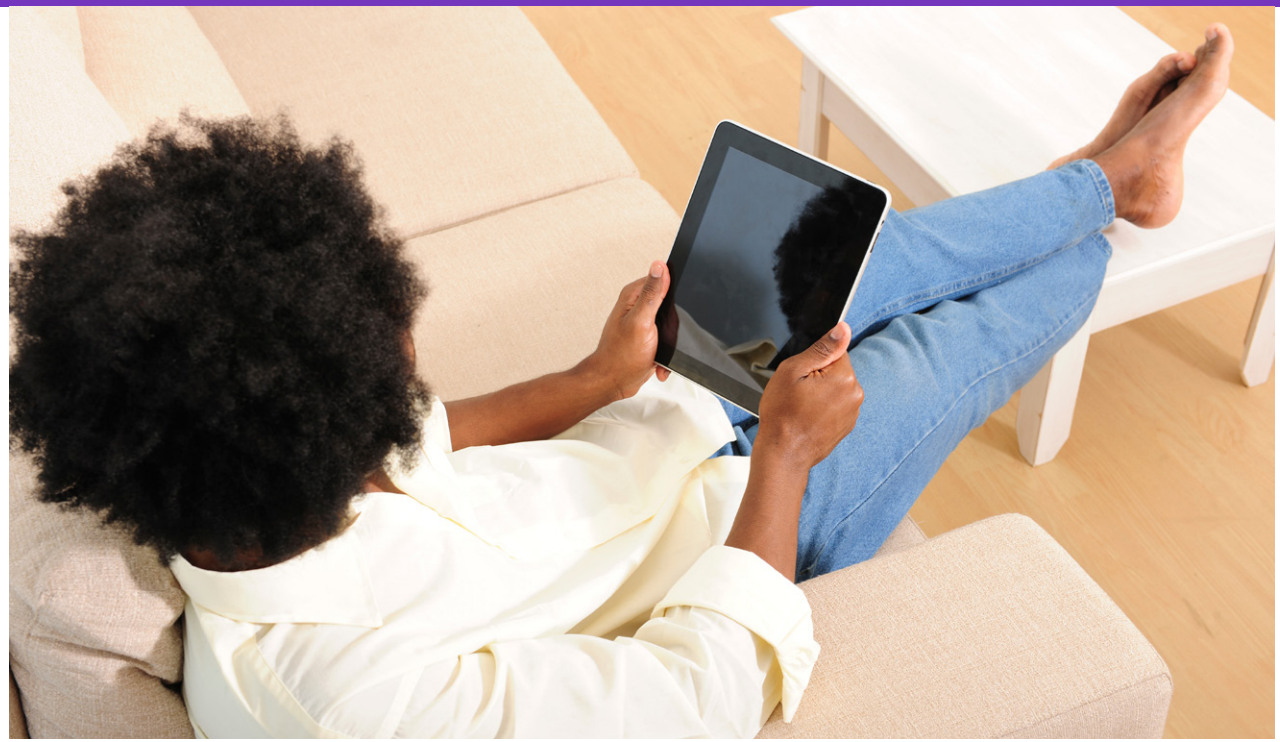

**Figure 2. Extending the content cloud throughout the home**

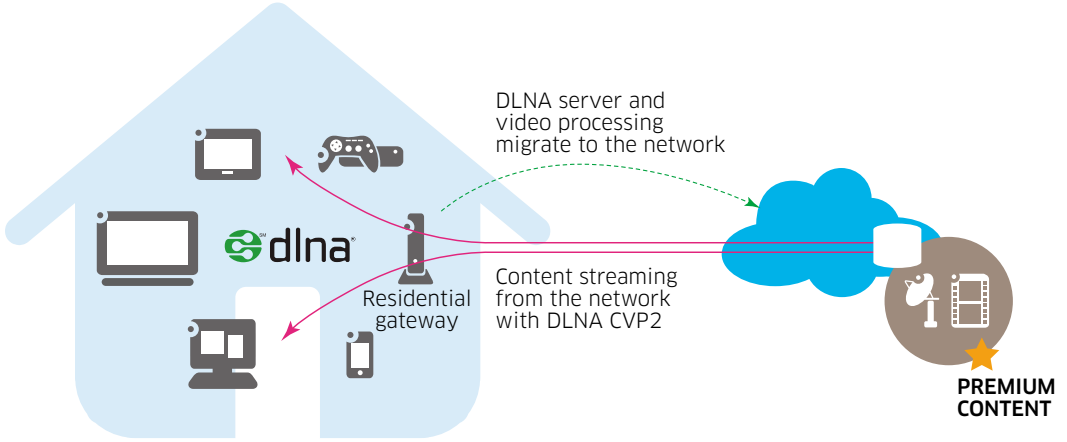

# Content Delivery Networks Get Personal

By Roland Mestric

**HIGHLIGHTS**

- Service providers can combine the CDN replication model with personalization
- Adding 2 key components increases CDN intelligence and processing power
- Manipulating content at the network edge increases scalability

## CDNS MAKE SENSE FOR PERSONALIZATION

Because **content delivery networks** (CDNs) are at the heart of next-generation IP video infrastructure, they are the logical place to add personalization. With the right approach, service providers can take advantage of the CDN replication model while delivering personalized content. At the same time, they can ensure the network scales to achieve quality expectations.

There are many good reasons to personalize content, including:

- **Targeted ads –** Advertisers may want to target ads at groups or individuals to increase their effectiveness. With targeted ads, end users watching the same TV show or video-on-demand receive different ads based on their profile.
- **Emergency alerts –** Local, regional, or national authorities may want to use TV messages to alert the public to an impending or current emergency in their area. Messages could include weather- or disaster-related notifications, health and safety bulletins, or missing persons alerts.
- **Blackouts –** Broadcasters may not want to air certain TV programs, such as major sports events, in a certain market. Content delivery may be blocked or substituted with alternate content for people living in the geographic area where the event is taking place. Managing blackout zones becomes more complex as content rights are fragmented across screens and as nomadic users get access to TV on their mobiles devices.
- **Quality adjustments –** Service providers may want to dynamically increase or reduce content quality based on contextual data such as subscriber profile, device used, and content watched. Adjusting content quality is a helpful tactic when bandwidth is scarce and when several end users are sharing the same bandwidth. Consider home networks and mobile cells where the access technology typically limits bandwidth to a few megabits per second. Service providers may want to control exactly how bandwidth is shared among the devices and give priority to premium subscribers and wide screens. Subscribers may also want to specify their preferences.

In all of these cases, end users receive different versions of the same content, or they receive alternate content for a limited time period. Most of the time, most of the content is common to all end users, allowing service providers to take full advantage of CDN efficiencies.

To use their CDN for personalized content, service providers need a way to partially or temporarily manipulate the content at the edge of their network. This means they must augment their CDN with capabilities that allow them to enhance HTTP generally and **HTTP adaptive streaming** (HAS) specifically.

## HAS IN A CDN ENVIRONMENT

The industry has adopted HAS for delivering multiscreen video to connected devices. HAS adapts the quality of the video stream in real time to match available bandwidth, helping service providers deliver video to more customers in a cost effective way.

One advantage of HAS is its compatibility with the standard HTTP caching mechanism. Service providers are building their own CDNs to replicate content at the edge of the network. The goal of the content replication mechanism for HTTP/TCP delivery is the same as that of multicast for Real Time Streaming Protocol (RTSP)/User Datagram Protocol (UDP) delivery in the IPTV world: They are both meant to reduce network traffic by serving individual end users from the edge of the network with a unicast stream. To achieve this goal, a CDN puts replication points in the network to copy the multicast tree.

Serving end users from the edge of the network allows service providers to **scale the network** and improve video quality for popular content.

**Figure 1. Like IPTV networks, CDNs serve end users from the edge of the network**

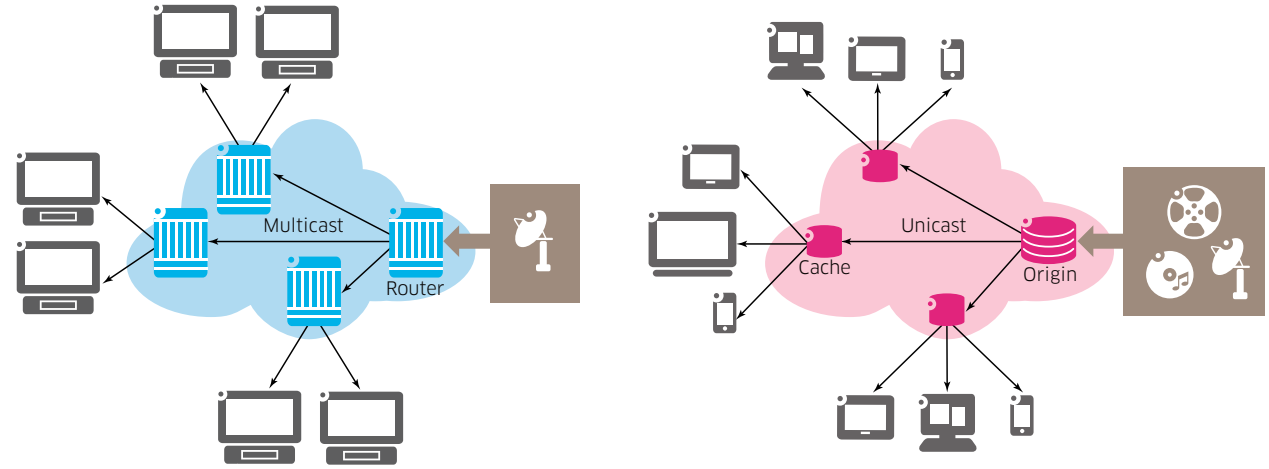


Service providers no longer need to send individual streams of the same content to every end user from a central point in the network. Instead content is sent only once across the network to a point as close as possible to the end user.

However, because HTTP is a stateless protocol, the streaming server does not manage the state of the client. The client is responsible for maintaining its own state. To the server, all client requests are equal, whether they come from the same client or multiple clients. As a result, service providers have little control over what is being delivered to the client. This makes it difficult for them to scale delivery cost effectively, and impossible for them to tailor content delivery for individual users or devices.

## NEW CDN COMPONENTS INCREASE CONTENT CONTROL

To take control of the delivery mechanism, service providers must build a server-side version of the client-based adaptive streaming concept. Introducing two new components will bring more intelligence and processing power into the CDN to improve overall HTTP behavior and performance:

1. A **session manager** retrieves the contextual information needed to customize the content and tell the cache in the CDN which changes to apply to the content.
2. A **video processor** in the cache generates the new content to be sent to end users based on the original content and the information delivered by the session manager.

### A session manager increases awareness

A session manager tracks each session and application. It helps service providers gain greater awareness of individual content requests and video streaming sessions. This awareness allows service providers to manipulate and personalize each video stream while accounting for current network conditions, subscriber profile, device type, and content being viewed. The data is collected from a variety of sources:

- Network conditions are retrieved from a real-time network analytics system.
- Subscriber profile and location are retrieved from a subscriber management system based on attributes such as device IDs and IP addresses.
- Device type is retrieved from HTTP GET requests sent by the devices.
- Content being viewed is retrieved from a content management system.

The session manager also interfaces with third-party systems that make decisions about changes that must be applied to the content. The type of system needed depends on the type of personalization required. Examples include:

- Ad campaign managers
- Emergency alert centers
- Broadcaster management systems
- Systems that support a Policy and Charging Rules Function (PCRF)

The session manager processes the information from the third-party system and passes on specific recommendations to the cache.

**Figure 2. A session manager and video processor increase control over content**

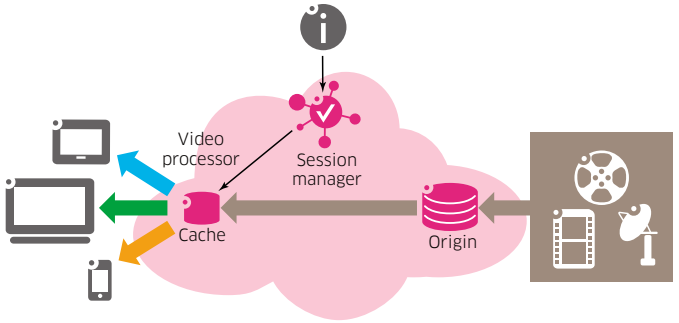

### A video processor customizes content

The video processor in the cache is a sophisticated engine with deep knowledge of the underlying video codecs and transport protocols. It uses the recommendations from the session manager to dynamically manipulate the original content for each end user or device.

The content manipulation performed by the video processor can include:

- Completely substituting the original content with other content during a limited time for a blackout or targeted ad.
- Slightly modifying the content by adding new flows to the original stream to display alerts or support trick modes such as pause, fast forward, and fast rewind as an overlay.
- Sending a different bit rate than requested or providing a limited list of available bit rates to control the quality delivered to end users.

## CONTENT MANIPULATION CREATES NEW OPPORTUNITIES

Content manipulation techniques also open the door to new opportunities to adapt, optimize, and secure video distribution to connected devices.

### Content protection

The ability to manipulate content at the edge of the network allows the CDN to better secure the content. Most content protection technologies rely on encryption techniques such as scrambling to restrict content usage to authorized users. Encryption can be applied at content preparation time, but its effectiveness is much higher if it is applied separately for each session. This technique is known as session-level encryption.

Inserting an invisible mark in the content before it is delivered to the end terminal is becoming a very valuable technique to protect content. It can be used as a complement or an alternative to classic digital rights management (DRM) solutions. The premium version of the invisible marking technology can personalize the watermark per video session. This makes it possible to track the end user who retrieved a specific asset from the network.

### Fast channel start

Unicast content delivery offers fast startup times that are comparable to IPTV with **fast channel change capabilities**. When an end user requests a video, the HAS client simply sends an HTTP GET request to the video server, and the server replies instantly with segments of the requested video. Video rendering begins as soon as a few video segments arrive in the HAS client buffer.

Unfortunately, important delivery functions such as entitlement checks, content manipulation, on-the-fly content encryption and decryption add delay to startup times. The delays are influenced by how responsive the systems providing these functions are, and by how well the end-to-end solution is designed, integrated and engineered. The additional delays can increase startup times by up to a couple of seconds or more. End users accustomed to the fast channel change capabilities that today's IPTV offerings provide will be disappointed by this performance.

Manipulating content in the CDN reduces startup time to improve end-user quality of experience. For example, the cache can serve the first few video segments as clear content to remove the delays associated with content encryption. Encryption is a time-consuming task. Bypassing this step for delivery of the first few video segments brings the benefits of the HAS fast startup time to the end-to-end content delivery solution.

## IMPROVE YOUR CDN EDGE

Service providers should treat the CDN as a key component of their core video network and not as an independent overlay. This approach exposes the CDN to intelligence held within the network. **Innovations from Alcatel-Lucent** enable the CDN to customize the content based on end-user preferences and context to deliver a better quality of experience.

The ability to manipulate content at the edge of the network allows any number of personalization possibilities while supporting scalable delivery. The same content is distributed cost efficiently to all end users across the CDN, and all content personalization occurs in the cache, close to the end user.

In contrast, solutions that personalize content in a central location dramatically increase traffic volume and storage requirements. Because they must support an individual unicast stream for each end-user request, centralized solutions are also harder to scale.

**To find out more about IP Video solutions, visit our website at www.alcatel-lucent.com/ip-video**



Alcatel·Lucent