VIDEO SHAKES UP THE IP EDGE

A BELL LABS STUDY ON RISING VIDEO DEMAND AND ITS IMPACT ON BROADBAND IP NETWORKS

STRATEGIC WHITE PAPER

As consumers demand ever more video and video consumption methods, the broadband IP networks used by many service providers to offer residential services will be overwhelmed by the sheer volume of data. The problem will be most acute at the IP edge of a triple-play provider's network, where video traffic will grow 2.5 times faster than traffic on subscriber links. To thrive in this new era of high bandwidth services, service providers will need to delight their customers with compelling new video on demand (VoD) services while keeping unicast video transport costs in check and controlling the bandwidth consumed by over-the-top (OTT) video.

To accomplish their goals, providers need to move away from broadband architectures based on legacy Broadband Network Gateway/Broadband Remote Access Server (BNG/BRAS) routers and adopt a distributed network architecture based on enhanced, terabit-speed BNGs and content delivery network (CDN) caches.

This document provides a detailed forecast of video consumption to 2020, determines the impact of video growth on IP service networks, and outlines strategies for rising above these challenges. A detailed Bell Labs case study quantifies the total cost of ownership (TCO) savings of moving from a centralized to a distributed IP edge with CDN caching.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

A growing consumer appetite for more video and video consumption methods will overwhelm the residential broadband networks many service providers use today. In the United States alone, Internet video consumption is expected to grow 12 times by 2020[1], and managed video on-demand (VoD) services are expected to grow 28 percent per year over the 5-year period until 2017. While the rising tide of video traffic will have an enormous impact on residential networks as a whole, the problem will be most acute at the IP edge of a triple-play provider's network, where video traffic will grow 2.5 times more than traffic on subscriber links.

To thrive in this new era of high-bandwidth video services, service providers must:

- Deliver the compelling new VoD services customers are asking for while keeping bandwidth and video delivery costs in check
- Control the escalating bandwidth and peering costs of over-the-top (OTT) video without sacrificing the freedom of choice subscribers crave

Accomplishing these goals requires a move away from broadband architectures based on legacy Broadband Network Gateway/Broadband Remote Access Server (BNG/BRAS) routers and centralized content injection. The operational costs of maintaining these networks will skyrocket from 2012 to 2017 as video traffic levels rise. Service quality will start to suffer as video servers and network links congest. The limited feature capabilities of traditional BNG/BRAS routers will limit a service provider's ability to offer personalized service experiences and control OTT traffic.

A popular strategy being adopted by an increasing number of service providers worldwide is to move to a distributed network architecture based on an enhanced, terabit-speed BNG and content delivery network (CDN) caches. A Bell Labs case study shows that a typical service provider can lower their total cost of ownership (TCO) by 33 percent from 2012 to 2017 with a distributed IP edge, with over 83 percent of savings coming from a reduction in bandwidth, peering and other operational costs. The TCO business case is strengthened by new revenue enabled by an on-net CDN, and the monetization of OTT traffic enabled through integrated application metering and control.

This document provides a detailed forecast of video consumption to 2020, determines the impact of video growth on IP service networks, and outlines strategies that allow service providers to thrive in this new era of high-bandwidth video services. A Bell Labs case study quantifies the TCO savings of moving from a centralized to a distributed IP edge with CDN caching.

[1] The source for all statistics in this paper is Alcatel-Lucent Bell Labs research, 2012.

VIDEO TRENDS AND FORECASTS

Bell Labs video forecasts were based on extensive research into the consumer trends driving online video consumption. The findings showed that consumers are looking for a more flexible video experience that enables:

- **Device freedom:** The ability to consume, control and share video content on any device — from PCs to gaming consoles, tablets and big-screen TVs.
- **Schedule freedom:** The ability to choose live TV, time-shifted TV or VoD.
- **Location freedom:** The ability to enjoy the "home" quality of experience (QoE) at public Wi-Fi® hotspots.
- **Content freedom:** The ability to watch any content type — cloud, premium, live or user-generated — at the resolution of their choice.
- **Interactive freedom:** The ability to integrate sidebar content, social media apps and interactive ads into the experience.
- **Plan freedom:** The ability to make tradeoffs between bandwidth usage and costs and not be subject to flat-rate, rigid caps.

To satisfy consumer demand for more choice and greater freedom, stakeholders in the video value chain — from content providers to tablet manufacturers to application developers — are investing in new video services, business models and technology innovations (such as the Apple® iPAD®) that are already making significant contributions to video consumption. This trend will continue well into 2020 and will be punctuated by:

- More video devices
  - ¬ Video on TVs, PCs, tablets, phones, gaming devices and personal digital assistants (PDAs)
  - ¬ Slim bendable screens for glasses and wristwatches; ultra-thin readers and video wall displays; integrated projectors
- More video business models
  - ¬ OTT video services
  - ¬ VoD in PayTV services
  - ¬ Increasing monetization through improved and personalized media search engines
- More video applications
  - ¬ Video insertion in cloud applications: gaming, business content, socializing/dating, education and health
  - ¬ Immersive video conferencing
  - ¬ Mass content creation and publishing by consumers
- More video quality
  - ¬ Increasing resolution and reality: high definition (HD), ultra-high definition (UHD), next-generation 3D and holographic imaging

The net effect of this activity will be to drive video consumption and video traffic levels in broadband networks to new highs on an almost continual basis. Bell Labs forecasts show dramatic growth over the next eight years (2012-2020). As shown in Figure 1, the total time spent watching video will grow from 4.8 hours to 7 hours per user per day. Much of this contribution will come from the latest generation of consumers, whose propensity for multi-tasking will result in 7 hours worth of video being consumed in as little as 5 hours. The proportion of time spent watching OTT and managed (PayTV) VoD services will grow from just over 33 percent to 77 percent. This will come at the expense of linear TV services, whose relative share of time will drop from 66 percent to 10 percent.

Figure 1. Time spent watching video, per user, per day in the United States in 2012 and 2020

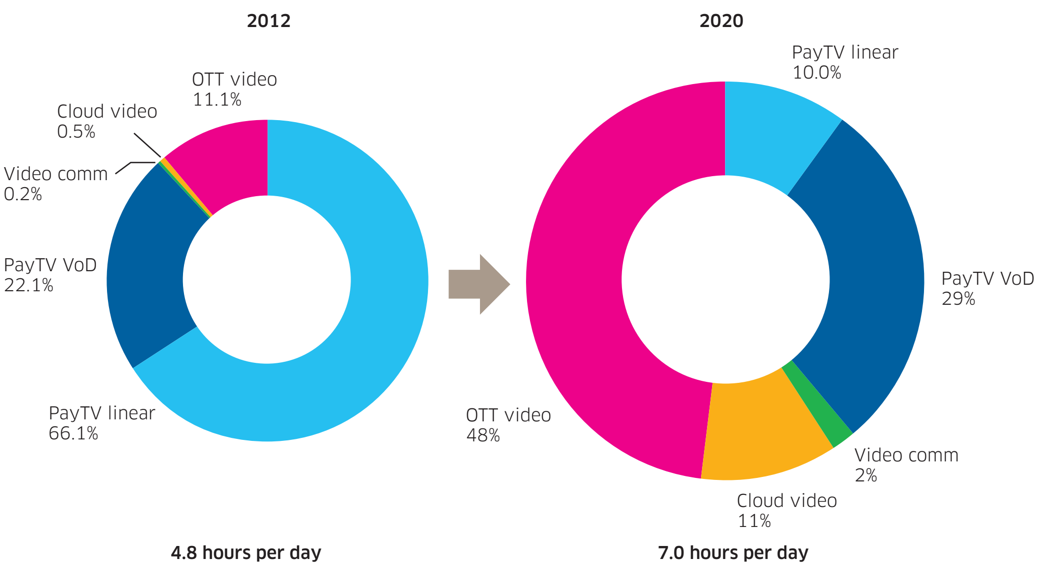

Source: Alcatel-Lucent Bell Labs, 2012

While users will spend almost half their time consuming OTT video, managed video services will continue to offer higher video quality and use more bandwidth per video streaming session. As shown in Figure 2, even though the blended play-rates of OTT video will increase twice as fast as the blended play-rate of managed (PayTV) video services, they will still lag behind by 30 percent at the end of 2020.

Figure 2. Blended play-rate forecasts for video applications in the United States

Source: Alcatel-Lucent Bell Labs, 2012

Combining the usage profiles in Figure 1 with the blended play-rates in Figure 2 gives an overview of video bandwidth demand from 2012 to 2020, which is shown in Figure 3. Internet video consumption is expected to grow 12 times, from 90 Exabytes to 1.1 Zettabyes. Growing user demand for VoD will also drive spectacular growth of managed VoD services, which are expected to show a cumulative annual growth rate (CAGR) of 28 percent, from 44 Exabytes to 244 Exabytes.

Figure 3. IPTV and Internet video demand forecasts in the United States, 2012-2020

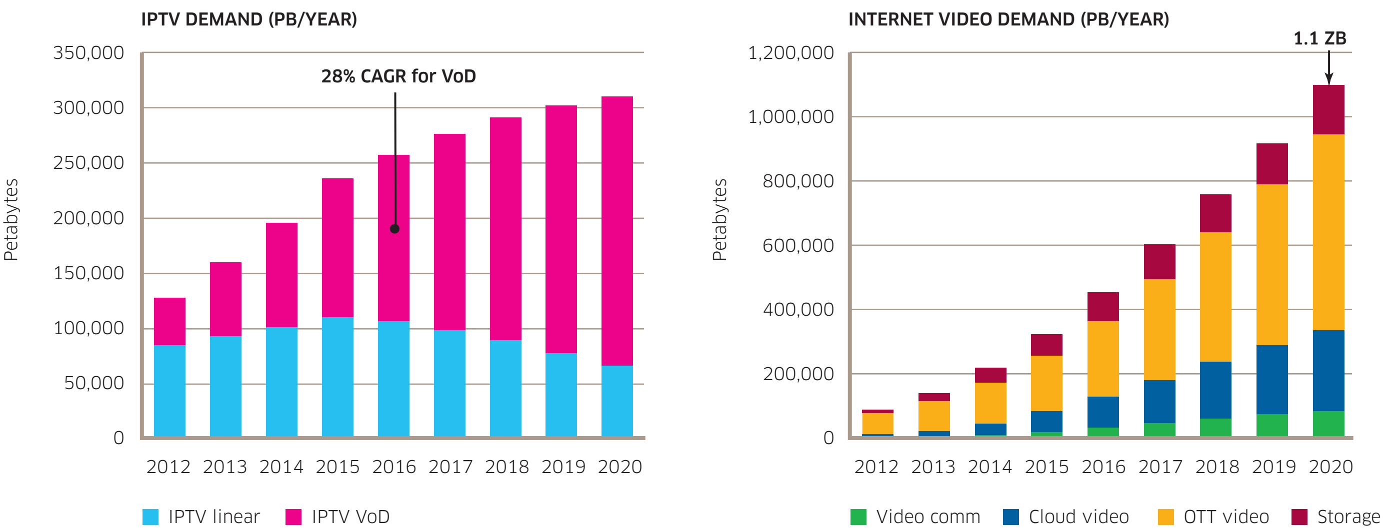

Source: Alcatel-Lucent Bell Labs, 2012

IMPACT ON THE IP SERVICE NETWORK

To determine how this increasing flow of content will impact a typical broadband network with a centralized BNG/BRAS architecture, Bell Labs first converted the video demand forecasts in Figure 3 into peak-hour bandwidth utilization forecasts — the point in time where broadband IP networks are most stressed. The peak-hour forecasts were then inserted into modeling tools that mirror and support the networks of Alcatel-Lucent's Tier 1 and Tier 2 broadband customers.

Figure 4 outlines how video applications and services are consumed during a 24-hour period. Some applications, such as cloud video, have a fairly even consumption profile during the day. Others, such as OTT and managed video services, are concentrated in the early evening hours. Just over 10.5 percent of managed video consumption and 8.5 percent of OTT video consumption will occur at the peak hour of 20:00 (8:00 p.m.).

Managed video service traffic is concentrated in a smaller number of providers compared to OTT video traffic. This fact, combined with higher blended play rates and higher peak-hour consumption, guarantees that in a typical triple-play service provider network, managed video traffic will continue to dominate the peak-hour bandwidth profile despite the rapid acceleration of OTT video.

Figure 4. Video demand forecasts in the United States, 2012-2020

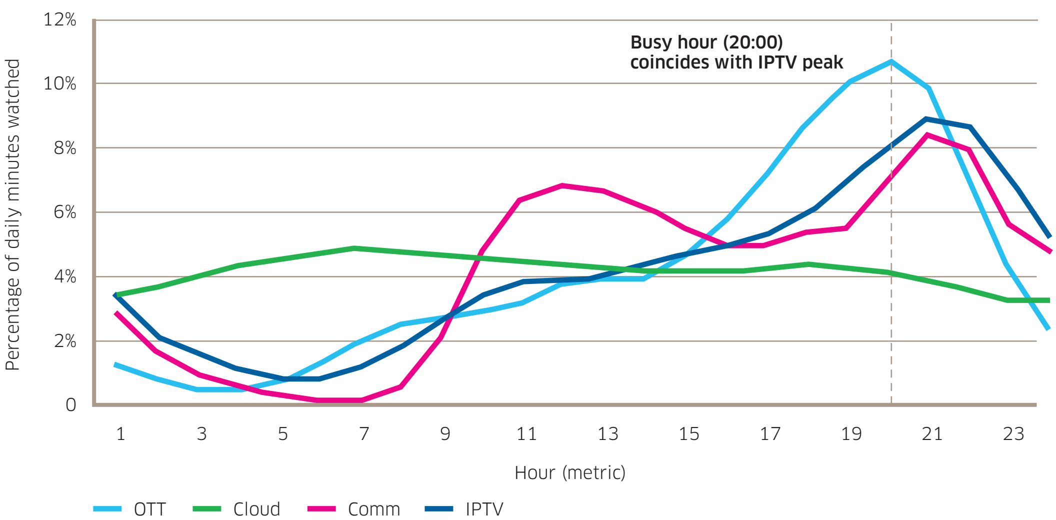

Source: Alcatel-Lucent Bell Labs, 2012

Impact analysis and service provider challenges

The peak-hour impact analysis results (see Figure 5) indicate that the video bandwidth crunch may happen much sooner than general video consumption forecasts suggest. Removing subscriber growth and all other factors from consideration, the data shows that video traffic at the IP edge of broadband networks will grow approximately 2.5 times faster than video traffic at the access layer and subscriber links.

Figure 5. Video traffic in a triple-play network during busy hour, 2012-2017

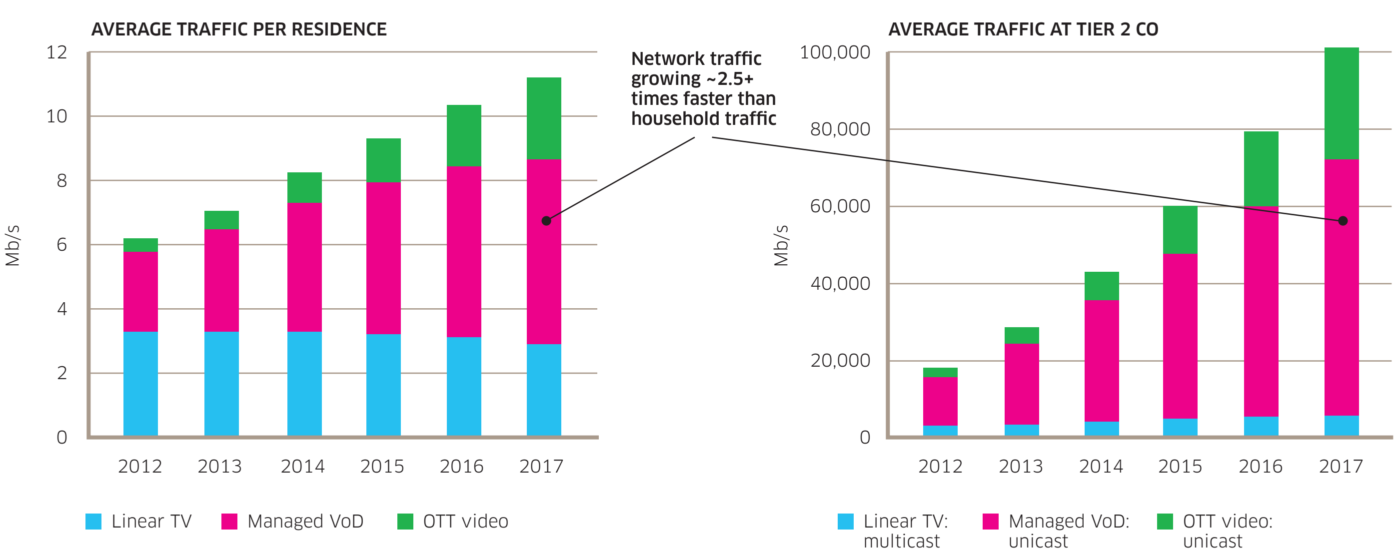

Source: Alcatel Lucent Bell Labs, 2012

This behavior can be explained by the shift from multicast to unicast streaming as subscribers opt for VoD services rather than linear TV. Whereas multicast delivery techniques leverage a single video server stream to reach hundreds of homes, unicast delivery techniques require a separate stream for each user, even when users are consuming the same content (see Figure 6). This scenario is played out in the real-world impact analysis shown in Figure 5. Linear TV bandwidth is high on subscriber links where all traffic is unicast, but is low in the IP service network where multicast efficiencies make their impact. The new VoD/unicast traffic represents a far larger share of overall traffic at the IP services edge than traffic at subscriber links.

Figure 6, Bandwidth overhead of unicast and multicast video streaming

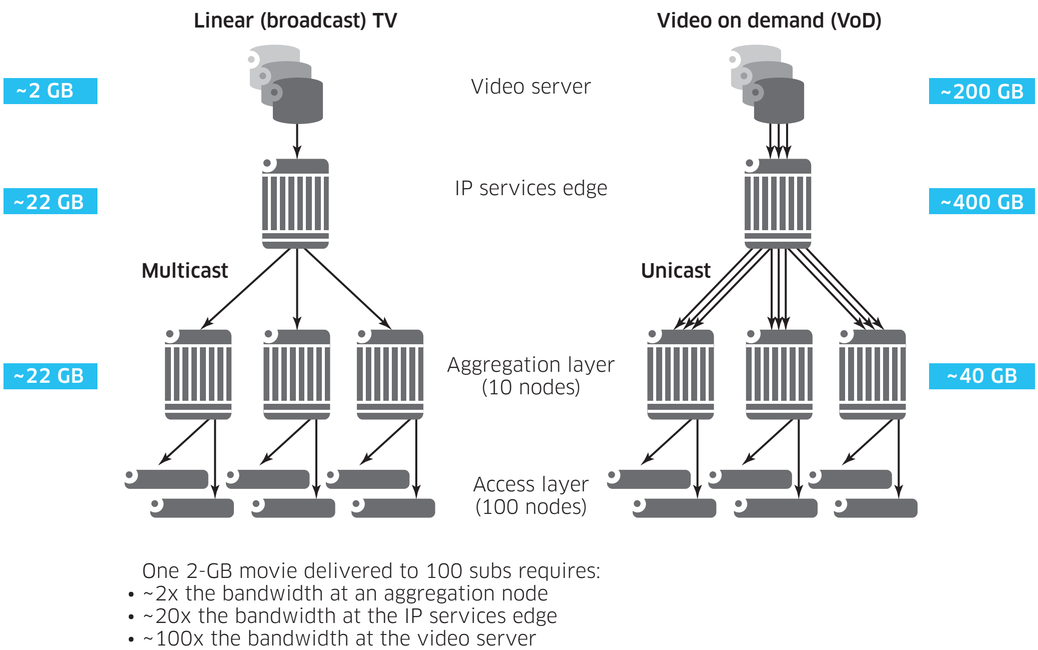

Source: Alcatel-Lucent Bell Labs, 2012

As video traffic levels in the IP service network accelerate, so do problems for service providers that rely on a centralized IP edge based on legacy BNG/BRAS routers:

- As unicast replaces multicast and video traffic levels surge, centralized content injection will drive up transport and peering costs in the IP service network, increase video server and link congestion, and decrease service availability and quality.
- The extra expense of maintaining the parallel networks shown in Figure 7, one for managed video and one for high-speed Internet (HSI) service, will become prohibitive. Parallel networks also isolate services from each other, making it difficult to deliver the blended service experience customers are asking for.
- BRAS routers simply do not have the capacity to handle the increased traffic load or the service intelligence to support the personalized video and Internet services customers are demanding.

Figure 7. Parallel broadband IP network architecture

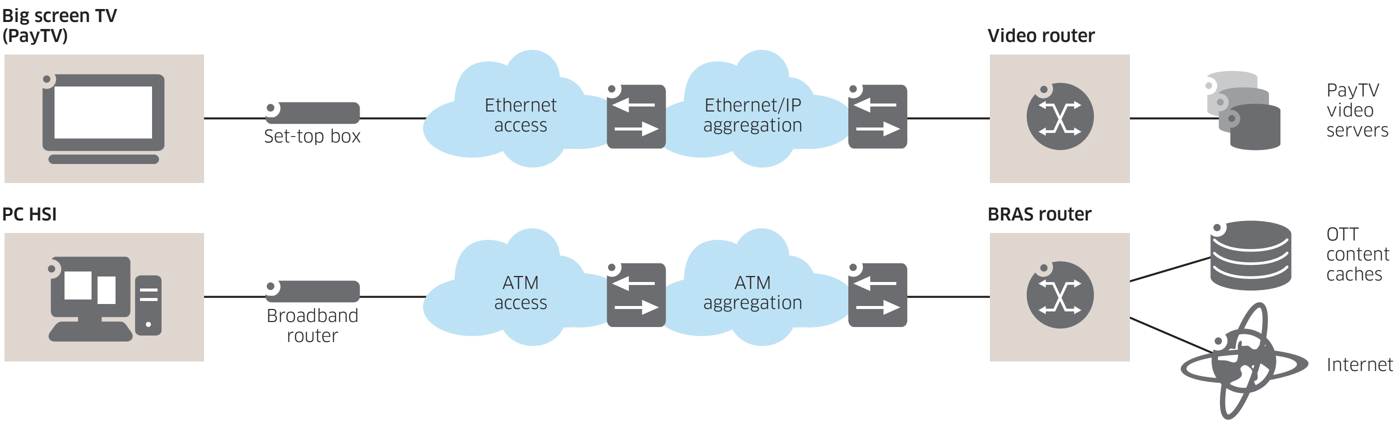

Source: Alcatel-Lucent Bell Labs, 2012

The preceding issues translate to two fundamental business challenges service providers must address if they are to embrace and monetize the surging tide of video:

- How to prepare the broadband IP network to deliver compelling new VoD services at the lowest cost per bit
- How to control escalating OTT traffic costs while delivering the freedom subscribers crave

PREPARING FOR THE VIDEO ERA: EVOLUTION STRATEGIES

Service providers wanting to thrive in the VoD era should consider investing in a distributed architecture that moves the BNG and the content injection point to the edge of an IP service network (see Figure 8). This new architecture includes:

- A terabit-speed, enhanced BNG that provides the capacity and performance to consolidate residential networks and accommodate surging video traffic levels for many years to come
- Integrated service intelligence to control and monetize OTT video traffic while delivering on subscriber demand for a more personalized broadband experience
- Distributed caches that increase service quality and reliability and enable video transport at the lowest cost per bit
- CDN software that enables optimal use of the distributed caches and a host of new video monetization opportunities as users move from linear TV to interactive, personalized VoD

Figure 8. Converged, next-generation broadband network architecture

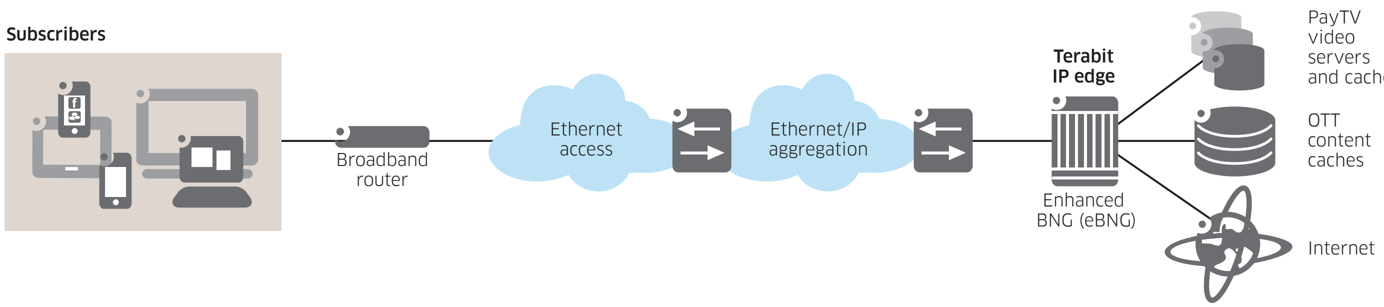

Terabit-speed enhanced BNG

To provide the performance levels and port density required to converge multiple residential networks and accommodate video growth for years to come, network processers at the heart of next-generation BNGs must be capable of 400-Gb/s speeds. All features must be supported at line rate so that network operators can create, monitor and control new feature-rich video services without impacting performance. Each chassis must support hundreds of thousands of subscribers with simultaneous access to multiple residential (managed VoD, managed linear TV and HSI) or business (Ethernet VPN, IP VPN and business Internet) services.

Integrated service intelligence

Consumers today are dissatisfied with current attempts by service providers to limit OTT VoD consumption, and often feel that HSI plans are designed to protect the service providers, not to benefit their customers. At the same time, the rigid use caps act as an obstacle to promoting the service providers (or their partners') VoD and related services.

What's required is a subscriber-friendly and more discriminating alternative — personalized plans that let the customer make the bandwidth trade-off decisions. This means integrated service intelligence within the BNG that can monitor traffic at the application level and limit, and boost and re-prioritize that traffic on a per-subscriber, per-session or per-application basis.

Before embarking on personalized plan offerings, service providers must first acquire detailed business intelligence on their subscribers' online habits. The enhanced BNG must be able to provide answers to the following questions:

Who are my top users?

What applications are they consuming?

How much bandwidth is required per application?

How do applications flow in my network?

Which applications use the most bandwidth, and when?

How do applications perform over time?

What percentage of traffic is OTT?

Armed with this information, service providers can begin offering personalized plans to ensure fair-share usage in a manner that satisfies each subscriber's unique needs. An example of a specialized plan is shown in Figure 9.

Figure 9. Example of personalized usage plans

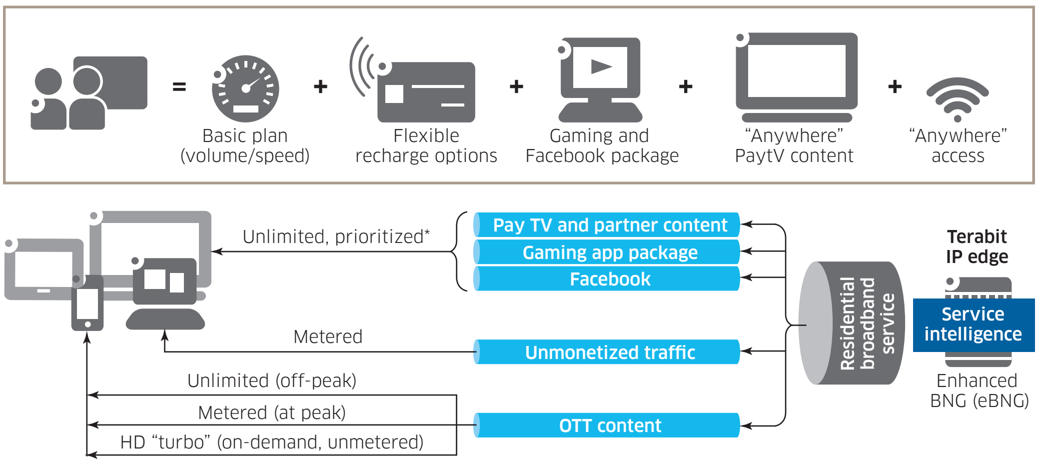

*Where applicable

Source: Alcatel-Lucent Bell Labs, 2012

Premium content is blended into the HSI service so that the subscriber can access it from any device in the home or through a community broadband link. All content-related traffic — including the content itself, the television guide and any related side-bar content — is identified by its application signature and is zero-rated so that it is not applied against the subscriber's usage cap. Selected applications, such as Facebook, are zero-rated for a fee or provided free as a promotional incentive.

OTT video applications are identified by their application signature and are metered. Unlimited OTT video viewing off-peak, or guaranteed high-quality boost for a period of time, can be used to minimize OTT traffic flows at peak times or to monetize the flows if bandwidth is available to support them.

Revenue-sharing arrangements can be negotiated to transform OTT video competitors into partners. The enhanced BNG can be used to identify, zero-rate or apply special treatment to partner OTT traffic flows as they traverse the IP network.

The enhanced BNG must also allow residential subscribers the freedom to take their tablets and other portable devices to Wi-Fi hotspots and connect as if they were at home — simply, securely, with access to all their service and content and with a similar QoE. Transparent Wi-Fi access allows residential broadband providers to offer subscriber freedom in the truest sense without overloading their 3G network.

Distributed IP edge and CDN caches

The introduction of a CDN and a distributed IP edge moves the video streaming origin point — and associated bandwidth load — from centralized video servers to distributed caches (see Figure 10) in Tier 1 and Tier 2 central offices (COs). The decision of which content to cache is based on dynamic demand, with stale content flushed as demand popularity declines over time.

The distributed caches are co-located or are integrated in the BNG, and are cost-optimized for delivery of popular content. Caches deeper in the network are cost-optimized to support a broader content library of less frequently watched content for a large audience.

The result is a significant bandwidth reduction between centralized video storage and distributed caches during peak times where concurrency is highest. And because the content travels a shorter distance to the subscriber, the QoE is far higher than anything an OTT provider can provide through their Internet-based streaming servers.

Figure 10. Distributed IP edge and CDN caches

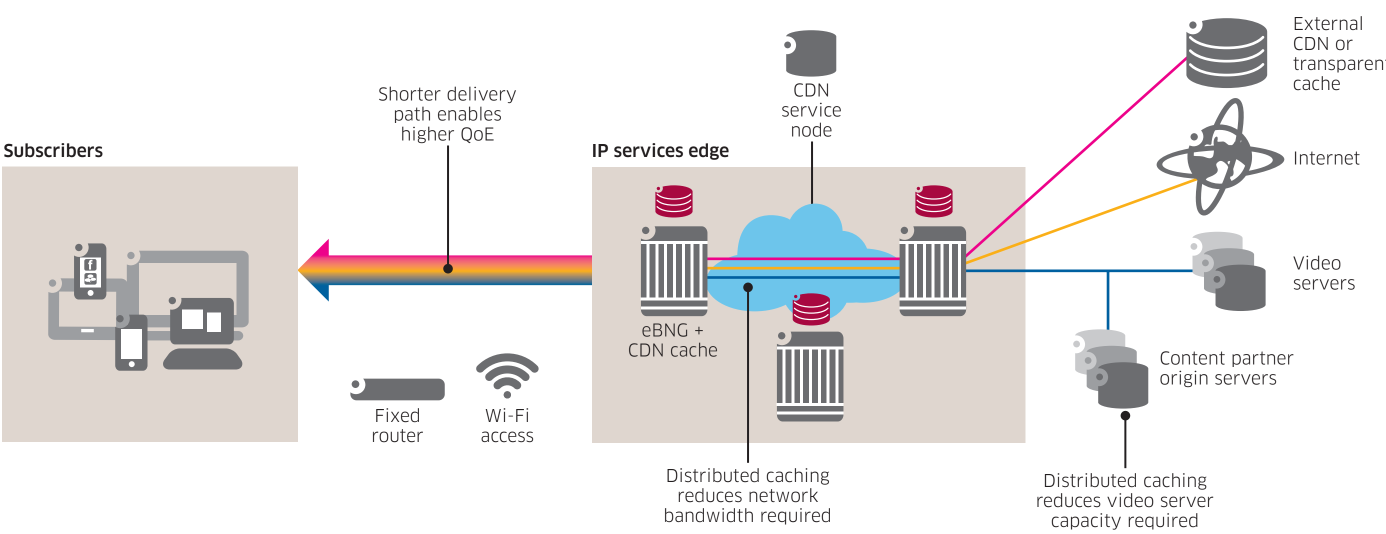

Source: Alcatel-Lucent Bell Labs, 2012

The CDN delivery model can be used in a complementary role to the linear TV delivery infrastructure to support both live broadcast and on-demand content requirements while meeting user expectations for quality, availability and reliability. Service providers can enhance their managed video services with broader content choices and evolve the traditional IPTV experience to an on-demand TV model as viewing habits evolve.

Service providers can also drive new revenue with CDN features that enable video service personalization, web interaction, multi-screen adaptation and personalized ad insertion. Content can be acquired by extending private content peering arrangements with a broad range of content providers or through content federations.

CASE STUDY: DISTRIBUTED IP EDGE AND CDN

This new broadband IP architecture is being deployed by top-tier service providers around the world as they replace their aging base of legacy BRAS/BNG routers, introduce new VoD services or transition their traditional IPTV networks. A Bell Labs case study compared the TCO of a centralized IP edge to a distributed IP edge with CDN content caches over the five-year period from 2012 to 2017. The network model was based on a large Tier 1 service provider in the United States and consisted of 1 video head office (VHO), 20 Tier 1 COs and 99 Tier 2 COs (see Figure 11).

Figure 11. Comparison of centralized IP edge and distributed IP edge

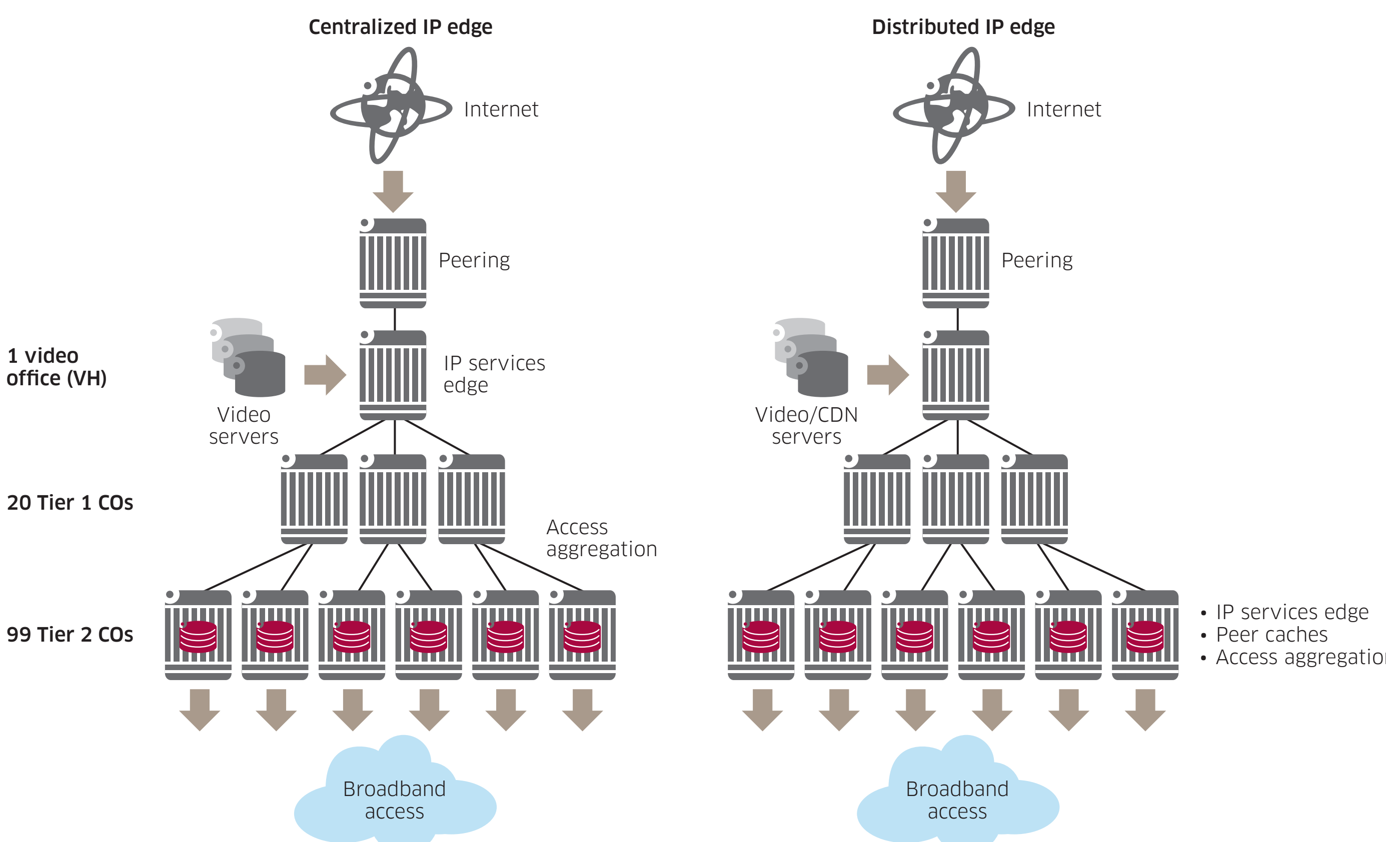

Source: Alcatel-Lucent Bell Labs, 2012

The study assumed all managed VoD traffic was served by the distributed CDN caches (see Table 1), and that the distributed cache hit rate would remain high (75-87 percent) due to the high concurrency associated with premium content. The distributed caches would also serve a small — but growing — portion of OTT traffic. The assumption is that OTT vendors would start to partner with service providers to increase the QoE associated with their service as a result of increasing pressure from application-level bandwidth caps.

The amount of OTT traffic addressed would remain relatively small for the first four years (only 15 percent) as would the cache hit rate due to the more diverse content library of an OTT provider.

Table 1. Cache hit rates for IPTV traffic and OTT traffic, 2011-2020

Year	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
CDN storage	4 TB					10 TB				
IPTV traffic addressed	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Cache hit rate	75%	75%	75%	75%	75%	87%	87%	87%	87%	87%
OTT traffic addressed	15%	15%	15%	15%	15%	40%	40%	40%	40%	40%
Cache hit rate	45%	45%	45%	45%	45%	71%	71%	71%	71%	71%

Source: Alcatel-Lucent Bell Labs, 2012

Operational cost assumptions were as follows:

- Transport: $3,200*/month to $4,000/month per 10G link
- Internet connection: $4,500/month per 10G port for peering; $2.50/Mb/s/month for transit
- Space: $5,905/rack/year
- Power: $0.11039/kW hour
- Maintenance: 10 percent of deployed equipment CAPEX

* All dollar amounts are in US dollars.

The results reveal a 33 percent TCO savings in moving from a centralized architecture to a distributed architecture (see Figure 12). Eighty-three percent of the overall TCO savings come from reductions in recurring operational costs. The positive business is further strengthened by the new revenue opportunities enabled by an on-net CDN.

Figure 12. TCO savings in move from centralized to distributed architecture

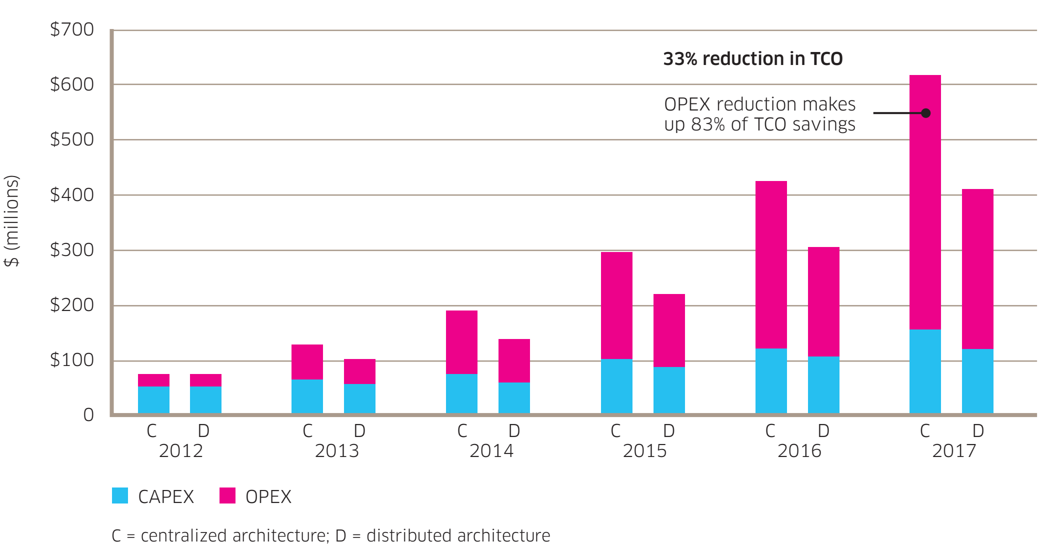

Source: Alcatel-Lucent Bell Labs, 2012

A closer look at operational costs (see Figure 13) reveals that operational costs are reduced by 36 percent in the distributed IP edge and CDN cache model. The bulk of savings coming from the 47 percent reduction in network bandwidth costs delivered by distributed caches as they reduce network transit bandwidth. Peering costs display a similar but less dramatic savings as OTT partnerships start to minimize the number of video streams originating from the Internet.

Figure 13. OPEX breakdown in move from centralized to distributed architecture

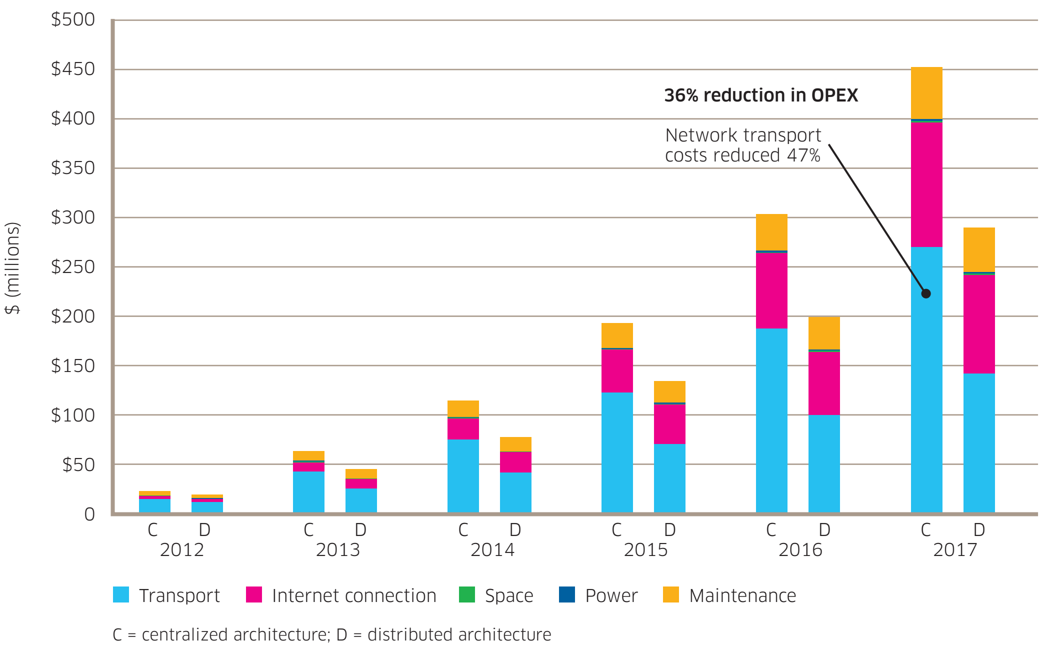

Source: Alcatel-Lucent Bell Labs, 2012

As shown in Figure 14, capital expenses were also reduced in the distributed model — by 33 percent. Although distributed caching was a new expense item in the distributed scenario, the investment was more than offset by a reduction in the number of video servers and network equipment required to support them.

Figure 14. CAPEX breakdown in move from centralized to distributed architecture

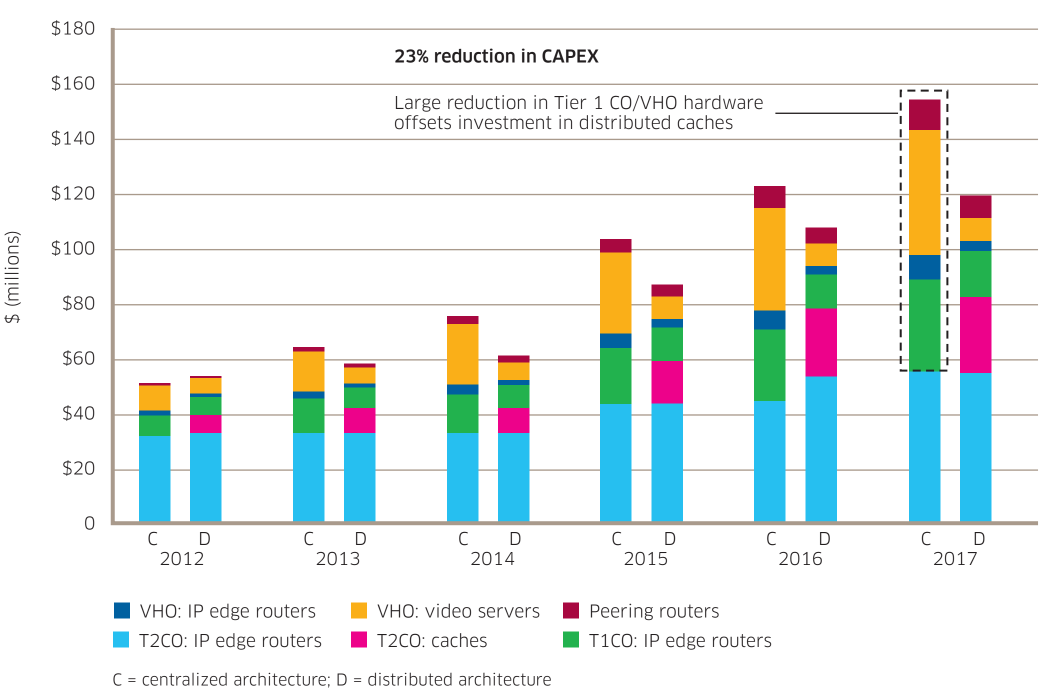

Source: Alcatel-Lucent Bell Labs, 2012

CONCLUSION

Video in all forms is driving revitalization in the IP edge of residential networks. Legacy BNG/BRAS routers deployed in a centralized architecture cannot scale to meet growing bandwidth demands, nor do they have the modern features required to deliver a more compelling and personalized residential broadband service experience.

By shifting to a distributed IP edge architecture — and employing modern high-performance, intelligent service routers and CDN caches — service providers can meet future residential bandwidth and services demand while realizing significant network savings.

Alcatel-Lucent is ideally positioned to help service providers through this necessary transition. Our extensive experience with residential service networks is anchored by the industry-leading service router platform, the Alcatel-Lucent 7750 Service Router (SR).

As a key element of the Alcatel-Lucent comprehensive residential network architecture, the 7750 SR is optimized for delivery of high-bandwidth, high-quality and highly flexible residential broadband services. It can be used as an enhanced BNG, which allows high-quality VoD, personalization and Wi-Fi access to be incorporated into the service provider's existing base of IPTV, HSI and voice services. These capabilities can play a key role in helping service providers capitalize on subscriber demands for a richer online experience.

ACRONYMS

BNG	Broadband Network Gateway
BRAS	Broadband Remote Access Server
CAGR	compound annual growth rate
CDN	content delivery network
HD	high definition
HSI	high-speed Internet
OTT	over-the-top
PDA	personal digital assistant
QoE	quality of experience
TCO	total cost of ownership
UHD	ultra-high definition
VoD	video on demand

www.alcatel-lucent.com Alcatel, Lucent, Alcatel-Lucent and the Alcatel-Lucent logo are trademarks of Alcatel-Lucent. All other trademarks are the property of their respective owners. The information presented is subject to change without notice. Alcatel-Lucent assumes no responsibility for inaccuracies contained herein. Copyright © 2012 Alcatel-Lucent. All rights reserved. M2012115052 (November)