Handling the Explosion of Online Video: Why Caching is the Key to Containing Costs

Author: Colin Dixon, Founder and Chief Analyst, nScreenMedia | Date: October 28, 2013

INTRODUCTION

Consumer thirst for online video seems to be unquenchable. With every passing month, more people join the ranks of online video viewers and another peak in usage is breached. According to comScore, on average, 190 million viewers each watch nearly 22 hours of online video per month.[1] Netflix subscribers watch even more, consuming almost an hour and a half per day.[2]

Fueling this dramatic rise in online video watching is the equally dramatic rise in the use of mobile devices, such as tablets and smartphones. Since the introduction of the iPad on January 27th 2010,[3] consumers simply couldn't buy them fast enough. In just three years, penetration of tablets in the U.S. has grown to 54% of online Americans and 44% of U.S. households with at least one of the devices.[4]

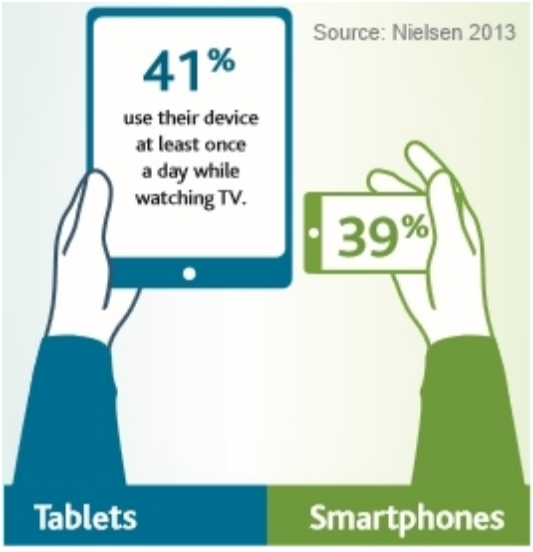

Figure 1. Mobile devices used while watching TV

The rise of mobile devices is also driving changes in the way consumers watch video. Over 40% of U.S. mobile device owners now watch television on a daily basis with a tablet or smartphone in hand. Consumers increasingly opt to watch their own shows rather than watching together on the big TV in the living room.

All of this video watching is having a dramatic impact on broadband networks. A recent Sandvine report found that Netflix video streaming now consumes 33% peak consumer bandwidth in the U.S.[5] As to total bandwidth consumed by all online video, Cox Communications reports that 65% of peak backbone bandwidth is now taken by over-the-top (OTT) video.[6]

The dramatic growth of video traffic on broadband networks is costing operators a lot of money. They must pay for bandwidth to interconnect their ISP network with the Internet, and must pay for expensive internal upgrades to cope with the extra video. Unfortunately, much of this bandwidth is wasted carrying exactly the same video to many different viewers. The latest movies, TV shows and viral videos are watched again and again by broadband users, yet each view of the video consumes another increment of video bandwidth.

This wasteful and expensive problem is increasingly drawing the attention of ISPs as they struggle to contain costs, maintain quality and scale their networks to handle the increasing load. One solution in particular looks like it may hold the key to solving this intractable problem: caching.

In this paper we discuss how caching can help broadband networks meet the ever increasing video demands of consumers. To put this topic in perspective, first we explain how much bandwidth is required to deliver video to primetime sized audiences, and how caching helps reduce the bandwidth required. Then we present partner caching and transparent caching along with the differences in their approaches to reducing the burden of video on operator networks. Finally, we examine the strengths and weaknesses of both approaches.

WORST CASE SCENARIO

We know where all the growth in online video watching is leading: every home will have at least one Internet HD video playing between 8 and 11pm, the peak TV viewing time. Unfortunately, it is likely worse than this. With our addiction to connected devices like iPads and smartphones, the average viewer will have one of these devices with them and will be watching videos while the TV is also streaming a show. In the worst case, each person in the household will be watching their own HD video stream on TV, and another one at the same time on a tablet or smartphone.

What does all this mean for bandwidth consumption? Let's find out.

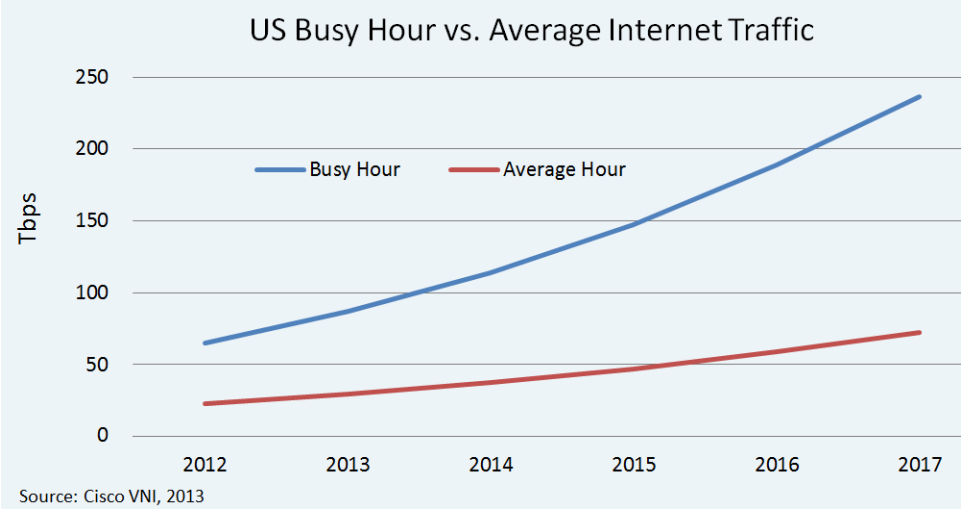

Figure 2. US Busy hour and average hour bandwidth requirements

Using an advanced video codec, like MPEG4 h.264, a very good HD picture can be delivered to a 50" TV with video streamed at 3.5 Mbps. Of course, if we were delivering sports this would probably be higher, and similarly a cartoon would probably be lower. With the smaller screens, delivering similar quality to a tablet and smartphone would require less bandwidth; I will assume 1.5 Mbps and 0.9 Mbps respectively.

If every one of the 116 million TV households in the US were streaming video to a connected TV,[7] that would require 347 terabits per second (Tbps) of bandwidth. If each of the homes were using a tablet, that would absorb 150 Tbps, and for a smartphone that would require 104 Tbps. If every U.S. TV viewer were watching TV and a connected device simultaneously, that would absorb 1.4 petabits per second. However, we are trying to model primetime and GfK says that 56% of people reported watching TV from 8 to 9 PM.[8] Based on this data, the worst case peak bandwidth estimate to support primetime using purely point-to-point, uncached video streams is 800 Tbps.

Is it probable that homes will have people watching video on all these different devices at once? This is already happening. Over 40% of owners of mobile devices use them while watching TV.[9] With live events like the Oscars providing mobile apps with live alternative camera views, it's likely this type of multiscreen simulcasting will be very common. At peak viewing times each of the 2.6 people per U.S. household may well be watching their own favorite show on one device with another stream going on their "second screen."

How does this compare with peak usage levels today? Cisco VNI forecasts that peak busy hour will consume 87 Tbps in the U.S. this year, of which video comprises a little over half. Video bandwidth required to support primetime on broadband is, therefore, almost 15 times greater than current requirements.

CACHING IS THE KEY TO SCALE

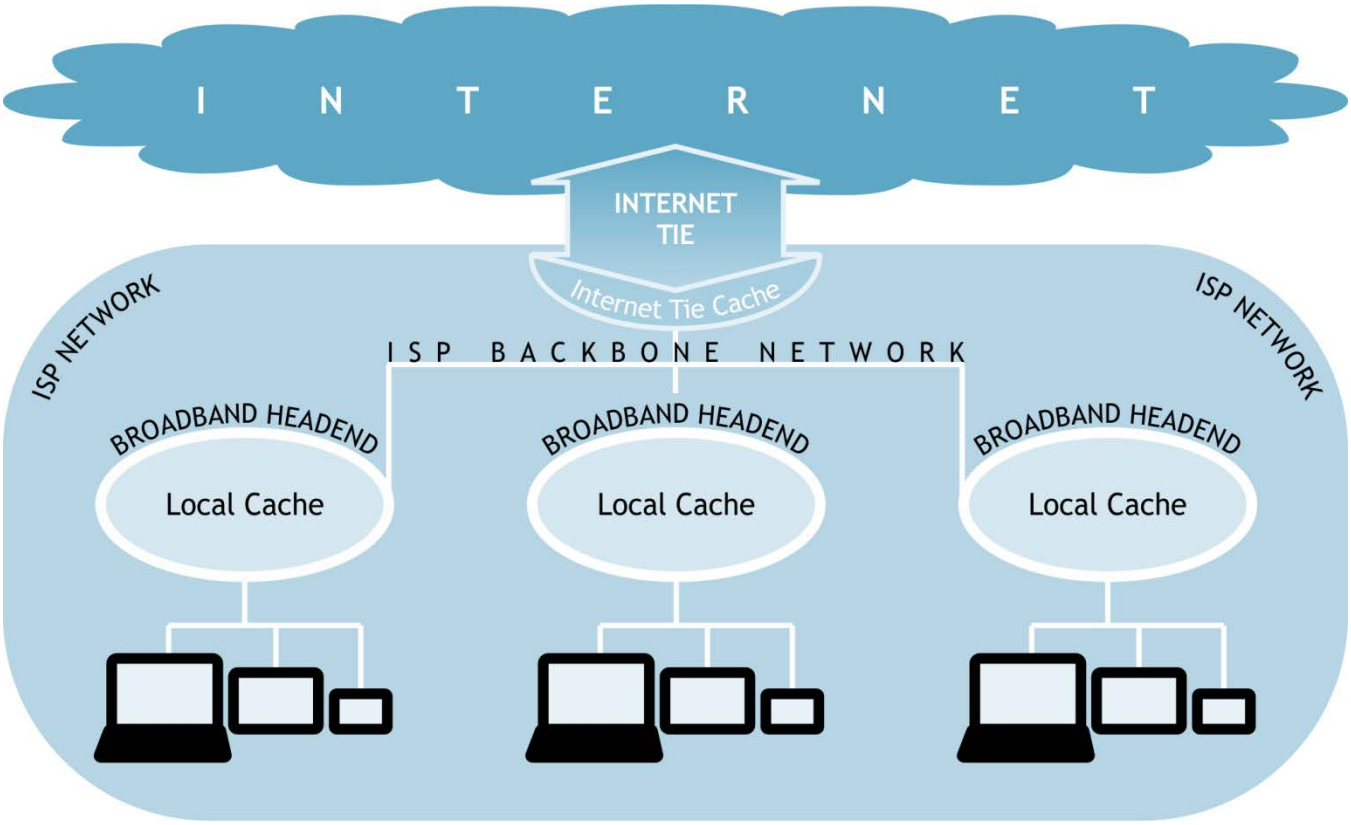

Figure 3. Places to cache in an ISP network to reduce Internet tie and core bandwidth traffic

Today the majority of video on the Internet is delivered point-to-point. That is, I request the video from a server in the cloud and it is streamed directly to me, and no one else. If my neighbor in California requests exactly the same video, he gets his own unique video stream. If the video we are watching is coming from Australia, our two streams, carrying exactly the same video, absorb twice the single stream bandwidth between Australia and California. Imagine what would happen if the 150 thousand [10] Australian expatriates living in North America all wanted to watch the AFL Grand Final simultaneously?

> The amount of bandwidth required to support primetime television over the Internet is 15 times greater than current requirements.

To prevent a disaster scenario such as this, the answer is to store copies of the videos people are watching close to where they live. When someone requests the video, it is streamed from the local copy. This technique is called caching and its impact on bandwidth consumption can be dramatic.

Consider the AFL Grand Final example above. To stream 150 thousand individual HD video streams of the game would require approximately 0.5 Tbps of bandwidth. All of these streams would be carried by the main connections between Australia and North America. However, if the video is cached locally in each of the 50 U.S. states, and in Mexico and Canada, just 52 streams, consuming 156 Mbps, would be required to be carried on the ties between Australia and North America. That reduces the bandwidth needs almost 3000 times.

Video caching is by no means a new idea. Streaming companies such as Akamai and Limelight Networks

use caching extensively in their networks. However, for ISPs the options to leverage caching technology are still relatively undeveloped. And this is costing them a lot of money.

For most ISPs, a major cost of the business is paying for the bandwidth that connects their broadband network to the Internet. The more bandwidth they use, the more they pay.[i] The cost of the interconnection today is typically $1.57 per Mbps. [11] To put this in perspective, if an ISP has 1 million subscribers, 20% of whom stream shows on Netflix simultaneously, the cost to cover that bandwidth is over one million dollars a month.

The bandwidth costs for the Internet tie, however, could pale into insignificance against the internal network costs. Upgrades to accommodate all the extra video bandwidth required within an ISP's network can be enormous. An upgrade to a single core router can cost tens thousands of dollars. [12] The ideal caching solution would not only reduce Internet tie bandwidth needs, but also internal network upgrade costs. To do both, there needs to be lots of local caches throughout an ISP's network.

Let's look at two options available to ISPs to help this problem: partner caching (p-cache) and transparent caching (t-cache).

Partner Cache (P-Cache)

Cox Communications found that as much as 50% of total traffic on the company's broadband network came from just 3 content providers. [13] What if Cox could partner with those 3 providers to put a cache of their content in Cox's network? The company could save as much as 50% of the bandwidth they currently pay for connecting to the Internet.

A partner cache is just that: a content provider positions its content on a streaming server in an ISP's network. When a consumer wants to watch a video from that provider, the stream comes from the p-cache in the ISP's network saving the stream bandwidth from crossing the ISP's Internet exchange connection.

> Partner Cache (p-cache): A content provider positions its content on a streaming server (the cache) in an ISP's network.

A p-cache, then, is very good at reducing the amount of traffic crossing the ISP's tie to the Internet. However, unless the partner is willing to provide lots of p-cache capacity distributed within an ISP's network, it won't help reduce bandwidth needs internal to the ISP's network.

Some content providers are very keen to work with ISPs to provide a p-cache. For example, Netflix has set up its own CDN (Content Distribution Network) called Open Connect. It offers to set up an Open Connect p-cache in an ISP's network for free. [14] Why would Netflix offer to install and maintain a cache for free in an ISP's network? Because it will save money doing so.

Netflix must pay for bandwidth to deliver its video through a traditional CDN like Akamai. If it can shift that traffic to Open Connect, the company stands to save a lot of money.

How does a partner cache like Open Connect work? Netflix's David Fullagar, Director of Content Delivery Architecture, says that the company uses a proactive caching technique where they preposition popular content in the cache. Doing this, he claims 100% of the content stored is actually accessed later on the network. [15]

On the face of it, the p-cache approach sounds like a win-win for ISPs and content providers alike. And ISPs like Frontier Telecom and CableVision have already signed up with Netflix Open Connect. [16] As we shall see, however, there are some disadvantages that might make it less interesting for an ISP.

i For large ISPs, it is a little more complicated than this as the amount of traffic delivered and received across the Internet tie offset each other. However, with the massive growth in Internet video consumption, even the large ISPs receive far more data than they deliver today.

Transparent Cache (T-Cache)

While installing a Netflix Open Connect p-cache can help a lot with the 33% of peak traffic the video service generates, it does nothing at all for the other 67% of traffic. Wouldn't it be nice if an ISP could install a cache that looks at all of the traffic coming into its network from the Internet, and cache all the most popular content locally? That would maximize the bandwidth savings and improve video performance for all customers. That is precisely what a transparent cache does.

Simply put, a t-cache looks at all the video[ii] on an ISP's network, identifies when something becomes popular, caches it on a local server and streams all further requests to watch the video from the local cache. From the point of view of the video provider, the operation of the t-cache is entirely transparent. Netflix, for example, would still see the user request to play, pause and fast-forward a video. The only difference is that the video is not delivered from the Netflix server, it comes from the local t-cache.

> Transparent Cache (t-cache): Examines all video on an ISP's network, identifies popular content, stores it locally and streams all further requests for it from the cache.

It is very important that the t-cache remain transparent to the streaming server. The streaming server needs to store statistics about playback and control who can watch the videos, all of which is critical to the provider's business. For example, counting video plays and ad views is important to ensure the video provider gets paid. The t-cache does not touch all of the messages between server and client concerned with these control functions. It only caches the video traffic.

Although an ISP must buy a t-cache solution, this gives it more freedom than with p-cache. An ISP can look at the cost-benefit trade-off of putting multiple t-caches in the network to reduce internal bandwidth. If going with a distributed t-cache solutions avoids the need for costly internal upgrades, it may well be cheaper to buy t-caches than to upgrade costly core routers.

T-cache solutions are also beginning to find favor with operators. Mediacom recently decided to employ Qwilt's QB series t-cache product for its network. JR Walden, VP of Technology at Mediacom, said the reason his company decided to go with a t-cache was because "we need a cost-effective way to extend our network infrastructure, while improving the quality of viewing experience." [17]

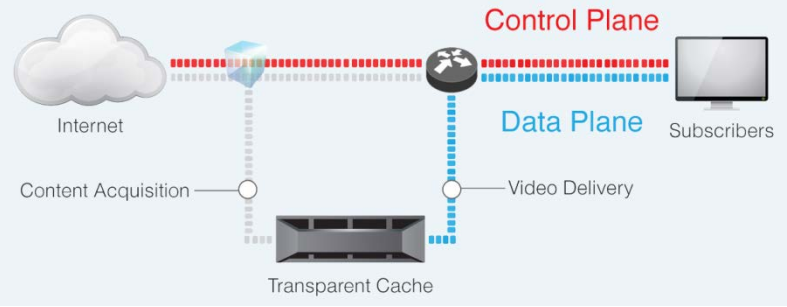

Figure 4. T-cache passes through video control traffic

[ii] T-cache works with other data types aside from video, including automated software updates and file downloads.

T-CACHE VERSUS P-CACHE

How do the two approaches to caching compare for operators? Let's look at the pros and cons of both in six different categories.

Equipment Cost

It is very hard to beat free, which is exactly what p-cache providers such as Netflix Open Connect and Akamai are offering. However, as we shall see, the potential operational costs of p-cache solutions, which operators must bear, could outweigh the upfront costs of t-cache in some cases.

Coverage of Video Traffic

Partner caches only store the content provided with the service. While this can be a boon for operators when one over-the-top video provider dominates, it becomes much less helpful the more providers there are. In the US market today, Netflix dominates video bandwidth usage, but over time this is unlikely to remain the case.

If we look at the dynamics of the video market, more and more online video providers are entering the online fray. At the same time, consumption of online video is increasing. Over the next several years, though Netflix will likely continue to grow, its share of overall bandwidth consumed will shrink. The future of online video looks to be headed to a much more fragmented market with many more providers.

In this situation, p-cache looks to be unmanageable for operators, who would be asked to manage an increasing number of different solutions. Transparent caching, which acts upon all traffic, is by far the best option to reduce bandwidth costs when many providers are involved.

New Operator Revenue

T-cache holds out the prospect of actually providing new revenue streams for an ISP. A CDN is contracted by a video provider to deliver video and is paid for the total amount of bandwidth delivered. For example, if a video provider is paying a CDN $0.20 a gigabyte to deliver video and customers stream 100 HD movies a day (about 200 GB of data,) the provider pays the CDN $40. If an ISP's t-cache delivers 20 of those movies, the video provider saves $4 in CDN costs.

How does an ISP get paid for doing some of the CDN's work? The t-cache tells the CDN it has delivered the videos so the CDN can bill the provider. In return, the CDN shares some of the revenue with the ISP. Although this model has still not been fully implemented, t-cache providers such as Qwilt and Akamai are close to finalizing the details. [18, 19]

P-cache provides no additional revenue for an ISP.

External and Internet Costs

As mentioned, there are two major costs of covering video bandwidth: Internet tie costs and internal network upgrades. Both p-cache and t-cache attack the external Internet tie costs, but only t-cache addresses the internal core network upgrade costs.

Impact on Quality of Experience

Both cache technologies should provide better video quality for online video viewers. Generally, eliminating as many sources of possible traffic congestion between a client and video server will improve the overall quality of the video delivered. Congestion negatively impacts video quality in two ways:

- Congestion on links causes video servers to lower the video quality delivered over that link. Lower quality video consumes less bandwidth
- Congestion and network failures contribute to video playback problems like slow video starts and pixilation of the image.

The penalties for these types of network problems can be severe. According to Conviva, a video provider will lose as much as 80% of its audience if the video takes 3 seconds or longer to start. Online video viewers are also becoming less tolerant of buffering (when a movie stops because video was not arriving fast enough to keep up). In 2011, a 1% increase in buffering resulted in 3 minutes less of viewing time per movie. In 2012, the identical 1% increase led to 8 minutes lost in viewing time per movie. [20]

> According to Conviva, if a video takes longer than 2 seconds to start, 80% of the viewers will leave.

If an ISP thinks this is just the video provider's problem, think again. A Netflix customer experiencing problems with streaming is very likely to call their ISP before they call Netflix.

Of course, t-cache should improve quality for all video on the network, while p-cache will only work for the specific service it is tied to.

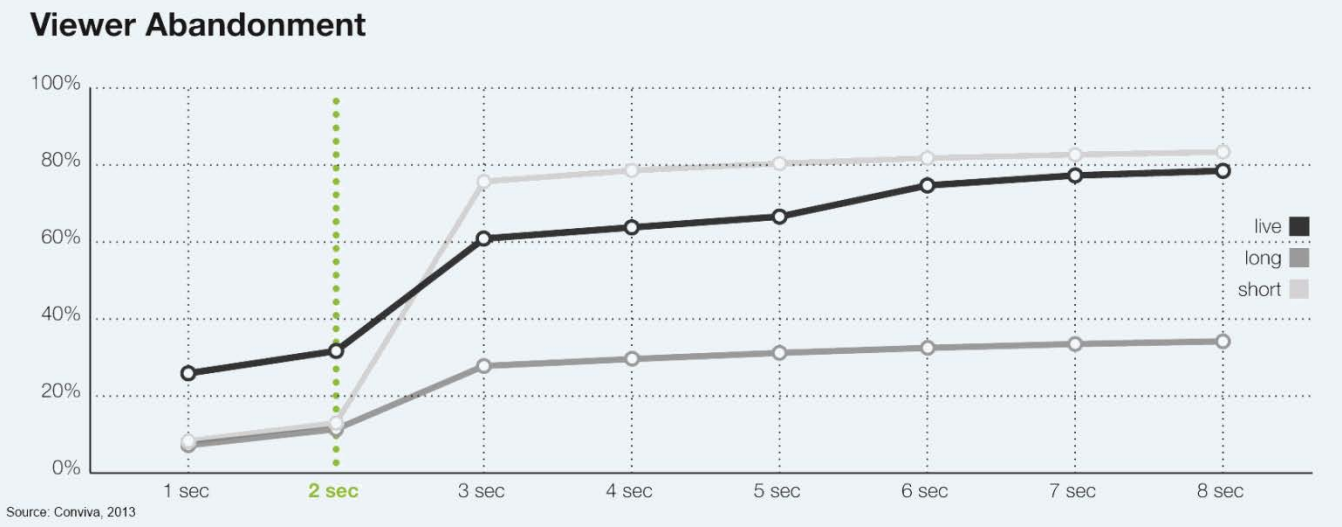

Figure 5. Most viewers won't wait more than 2 seconds for a video to start playing

Network complexity

We have already discussed that the trajectory of web video is headed toward many providers. A lot of these will want to optimize the delivery of video on an operator's network by providing a p-cache. This could result in a huge, unwieldy and expensive network to maintain. By contrast, a single t-cache solution should be able to handle the explosive growth in video providers, and in the bandwidth they will require.

Table 1. The pros and cons of t-cache and p-cache

	T-CACHE	P-CACHE
Equipment Cost borne by…	✘ Operator	✓ Video provider
Video traffic eligible for cache	✓ All	✚ Video provider only
External & Internal cost reduction	✓ Yes	✚ External only
New operator revenue	✓ Yes	✘ No
Quality of Experience improved	✓ Yes, all video	✚ Yes, provider only
ISP network complexity	✓ Less	✚ More

CONCLUSION

Caching is an essential part of the online distribution of video. Without it, point-to-point delivery guarantees huge amounts of expensive bandwidth tied up in duplicate movie and show streaming.

Employing P and T cache solutions can dramatically impact the bandwidth required to deliver online video over broadband networks and help restrain the associated costs.

> Caching is an essential part of the online distribution of video. Without it, broadband simply won't scale to primetime audiences.

P-cache solutions are very useful when there are a few very large online video providers. Since the partner bears the cost of the equipment, it is a cheap way to get instant savings on Internet tie bandwidth needs.

In the case where there are a larger number of online video providers with no single provider totally dominating, p-cache becomes a lot less attractive. An operator would need to support many p-cache solutions to yield useful reductions in bandwidth consumption. The savings in equipment costs from free p-cache systems will likely be outweighed by the increase in network complexity and operational costs.

T-cache solutions are excellent for when there are a large number of online video providers with no single provider dominating. Acting across all video traffic on the network allows t-cache solutions to maximize bandwidth savings on Internet ties and ISP backbones. The prospect of additional revenue through CDN partnerships, and reduced network complexity (in comparison to p-cache) should quickly erase the cost of purchasing the equipment.

In the case where there are a few dominant online video providers, t-cache still provides substantial bandwidth savings, but the initial equipment cost may make free p-cache solutions (if available from the provider) a more cost-effective approach.

Finally, regardless of the number of online video providers, if an ISP has internal network bandwidth problems, a distributed t-cache solution may be a viable alternative to costly internal network upgrades.

References

1 comScore, "comScore Releases August 2013 U.S. Online Video Rankings," *comScore Media Metrix*, September 2013, www.comscore.com/Insights/Press_Releases/2013/9/comScore_Releases_August_2013_U.S._Online_Video_Rankings (accessed on 9/28/13)

2 Ben Fritz, "Netflix streams more than 1 billion hours in June," *Los Angeles Times*, July 3, 2013, articles.latimes.com/2012/jul/03/entertainment/la-et-ct-netflix-billion-hours-20120703 (accessed on 9/28/13)

3 Apple Press Release, "Apple Launches iPad," January 27, 2010, www.apple.com/pr/library/2010/01/27Apple-Launches-iPad.html (accessed on 9/28/13)

4 Cellular-news, "Tablet Penetration in the U.S. Grows by 47% Since 2012," *cellular-news*, June 10, 2013, www.cellular-news.com/story/60419.php (accessed on 9/29/13)

5 Don Reisinger, "Netflix gobbles a third of peak Internet traffic in North America," *cnet*, November 7, 2012, news.cnet.com/8301-1023_3-57546405-93/netflix-gobbles-a-third-of-peak-internet-traffic-in-north-america (accessed on 9/28/13)

6 Jeff Finkelstein, speaking at CDN World in Las Vegas, July 2013

7 Nielsen, "Nielsen Estimates 115.6 Million TV Homes in the U.S., up 1.2%," *Nielsen NewsWire*, May 7, 2013, www.nielsen.com/us/en/newswire/2013/nielsen-estimates-115-6-million-tv-homes-in-the-u-s---up-1-2-.html (accessed on 8/26/13)

8 GfK, "Primetime TV 2004-20012 Among Persons 18-49," GfK White Paper, October 2012, marketing.gfkamerica.com/101212-1269/Primetime_TV_2004-2012_A_GfK%20Whitepaper.pdf (accessed on 8/26/13)

9 Nielsen, "State of the Media: Cross-Platform Report, Q2 2012," *Nielsen Reports*, November 13, 2012, www.nielsen.com/us/en/reports/2012/state-of-the-media--cross-platform-report-q2-2012.html (accessed on 10/3/13)

10 Southern Cross Group, "Estimated number of Australians Living Oversees," *Southern Cross Group*, December 2000, www.southern-cross-group.org/archives/Statistics/Numbers_of_Australians_Overseas_in_2000_by_Region_14_April_2001.pdf (accessed on 8/27/13)

11 Dr. Peering, "Internet Transit Prices — Historical and Projected," *Dr. Peering*, October 3, 2013, drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php (accessed on 10/3/13)

12 ISP Trader, "Cisco 7206VXR High Power Routers," *ISP Trader*, www.isptrader.com/category/product-lines/routers/aggregation-routers/cisco-7200-vxr-and-ubr (accessed on 10/4/13)

13 Jeff Finkelstein, speaking at CDN World in Las Vegas, July 2013

14 Ryan Lawler, "In The Fight Between Netflix And Cable Operators, High-Quality Streaming Video Being Held Hostage," *techcrunch.com*, January 16, 2013, techcrunch.com/2013/01/16/netflix-cable-blah-blah-blah (accessed on 8/28/13)

15 David Fullagar, Netflix, speaking at CDN World in Las Vegas, July 2013

16 Netflix PR, "Netflix Open Connect Delivery Network," *Netflix*, signup.netflix.com/openconnect (accessed on 9/29/13)

17 Qwilt PR, "Mediacom selects Qwilt for Online Video Delivery and Improved Quality of Experience for Subscribers," *Qwilt*, September 25, 2013, www.qwilt.com/news-and-events/2013/09/25/mediacom-selects-qwilt-online-video-delivery-qoe (accessed on 9/29/13)

18 Frank Childs, Director Product Marketing, Aura Network Solutions, Akamai, conversation with the author, Sept 27, 2013

19 Dan Sahar, VP & Co-founder of Qwilt, conversation with the author, July 19, 2013

20 Conviva, "Viewer Experience Report," *Conviva*, Q1 2013

www.nScreenMedia.com
For more information contact:
Info@nScreenMedia.com

This paper is made possible by the generous contribution of:

Qwilt

www.Qwilt.com