White Paper

Making the Case for Edge Data Centers®

Prepared by

Alan Breznick
Contributing Analyst, Heavy Reading
www.heavyreading.com

on behalf of

www.edgeconnex.com

January 2017

Introduction: Battling the Traffic Bottlenecks

With more people using the Internet, broadband speeds rising, more consumers watching online video, the number of mobile devices exploding and webpage sizes steadily increasing, Internet traffic is exploding and shows no signs of slowing any time soon. Consequently, cable operators, telcos and other communications service providers (CSPs) find themselves in a growing bind these days. As Web traffic keeps surging by as much as 30 percent to 50 percent a year, CSPs are seeing their backbone networks increasingly filling beyond the breaking point, clogged with the mounting cascade of data, voice and especially video signals as their customers keep consuming more video and adopting new connected devices.

At the same time, cable operators and other broadband providers are seeing their total network transport costs soar ever higher, even though the cost per port or data packet continues to decline. Despite their best efforts to boost efficiency, the overall increase in traffic carried over their networks is easily outpacing the unit decreases that they can manage to achieve through greater network efficiencies, as shown in **Figure 1**. Yet, if they want to keep meeting the growing customer expectations for the important services they are providing, operators have no choice but to carry the constantly heavier traffic load.

Figure 1: Residential Subscriber Busy Period Bandwidth Requirement by Application

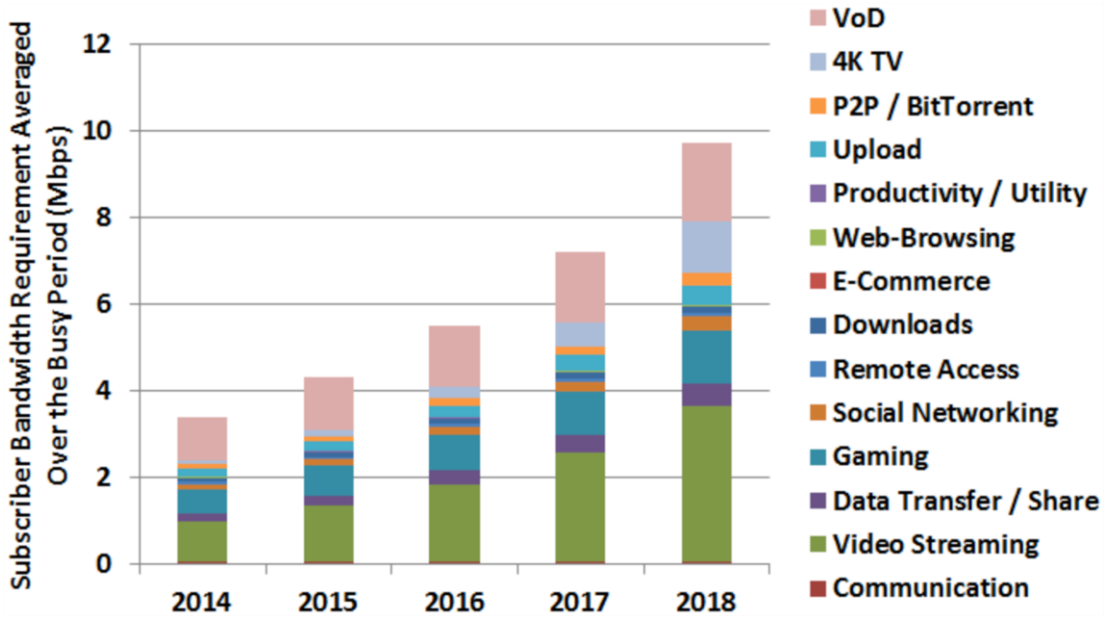

Source: ACG Research

As if these technical and cost issues are not imposing enough, broadband service providers are facing fresh challenges on the regulatory front, too. Thanks to the strong new net neutrality rules adopted by the Federal Communications Commission (FCC) last year, cable and other providers must take great care not to favor one piece of content over another or prioritize traffic for one user over another on their networks. As a result, CSPs are limited in their ability to use some of their traditional network management tools and procedures to ease the growing traffic congestion.

Consumers are feeling the impact of these trends, as well. As the traffic bottlenecks escalate on the providers' backbone networks, cable and other subscribers are seeing slower data service, more video buffering, more frequent and bigger service interruptions, and other mounting problems. Subscribers are also experiencing more frequent dropped connections, greater latency and other service delays caused by the traffic growth as they try to add tablets, smartphones and other new mobile devices to their already strained home networks.

Given all these developments, what's a beleaguered cable operator or other CSP to do? Fortunately, there are some steps that providers can take to improve their network management, boost their performance and deliver a better experience to their subscribers without increasing their capital and operational costs further or running afoul of the uncertain and potentially volatile regulatory environment. They do not have to settle for the uncomfortable and increasingly expensive *status quo*. Instead, they can tap into ways to leverage infrastructure to increase performance and enhance user experience while lowering costs.

This white paper examines the latest trends in the broadband and video space and examines their implications for cable operators and other CSPs. Then it explores what providers can do to overcome the challenges presented by the growing traffic bottlenecks.

Specifically, this paper focuses on the use of localized Edge Data Centers® (EDCs), rather than the standard large, centralized data centers in far-off locations, to deliver bandwidth-intensive content and latency-sensitive applications closer to end users. It spells out what EDCs are, describes how they work and compares their performance with larger, centralized data centers. It shows how EDCs can produce significant transport cost savings for CSPs both nationally and regionally, while also improving the experience for their customers. Finally, the paper cites real-world examples to show just what can be achieved.

The Growing Traffic Congestion Problem

As one study after another has made clear over the last couple of years, Internet traffic is growing by leaps and bounds right now, with no end seemingly in sight. In fact, various estimates see Web traffic surging by as much as 30 percent to 50 percent annually, depending upon the market segment chosen. Moreover, Internet traffic is expected to continue surging through at least the end of the decade, as the Internet becomes ever more central to the life and well-being of the entire planet.

For instance, Cisco Systems' latest annual Visual Networking Index (VNI) forecast projects that global Internet traffic will grow threefold from 2015 to 2020, producing a compound annual growth rate (CAGR) of 25 percent over that span. As a result, its most recent forecast calls for Internet traffic to reach an astounding 161.3 exabytes per month in 2020, up from 53.2 exabytes per month in 2015.

Even looking just at North America, the traffic numbers are still quite impressive. The latest VNI forecast predicts that North American Internet traffic will grow almost threefold from 2015 to 2020, generating a compound annual growth rate of 23 percent. That means that traffic will soar to 44.7 exabytes per month by 2020, up from 15.8 exabytes per month in 2015.

There are numerous reasons for this ongoing surge in Web traffic. First of all, the number of Internet users around the world is still climbing rapidly. In its most recent VNI

forecast, Cisco projects that there will be nearly 4.1 billion Internet users globally by 2020, up from 3 billion users last year. In other words, more than 52 percent of the world's entire population will be using the Web by then.

Second, the number of devices that people use to access the Web continues to multiply, as does the number of machines accessing the Web, including Internet of Things (IoT) devices such as refrigerators and cars. Citing the latest Cisco VNI forecast again, the number of networked devices and connections is projected to climb to a whopping 26.3 billion by 2020, up from 16.3 billion in 2015. That would amount to an average of more than three devices for every man, woman and child on the planet, up from about two devices per capita now.

The inexorable rise in broadband speeds is another major factor contributing to the Internet traffic growth. With numerous CSPs around the world now starting to offer to their broadband subscribers data download speeds of 1 Gbit/s or higher, as well as faster tiers below 1 Gbit/s, subscribers are receiving and using more network capacity than ever before, with no end in sight once again. Indeed, the latest Cisco VNI forecast projects that the average fixed broadband connection speed will nearly double globally over the next five years, climbing from 24.7 Mbit/s in 2015 to 47.7 Mbit/s by 2020.

Even more notably, the continuing shift to IP video is driving much of the growth of Internet traffic. With the rise of such new video entertainment services as streaming video, subscription video-on-demand, over-the-top video and now, 4K Ultra HD (UHD) video, more and more bandwidth is becoming dedicated to video transport, even though the Internet was never meant for that.

Figure 2: The Big Shift to IP Video

OTT video accounts for more than 70% of peak traffic, but cloud & streaming are growing daily

Upstream		Downstream		Aggregate	
BitTorrent	18.37%	Netflix	35.15%	Netflix	32.72%
YouTube	13.13%	YouTube	17.53%	YouTube	17.31%
Netflix	10.33%	Amazon Video	4.26%	HTTP – Other	4.14%
SSL – Other	8.55%	HTTP – Other	4.19%	Amazon Video	3.96%
Google Cloud	6.98%	iTunes	2.91%	SSL – Other	3.12%
iCloud	5.98%	Hulu	2.68%	BitTorrent	2.85%
HTTP – Other	3.70%	SSL – Other	2.53%	iTunes	2.67%
Facebook	3.04%	Xbox One Games Download	2.18%	Hulu	2.47%
FaceTime	2.50%	Facebook	1.89%	Xbox One Games Download	2.15%
Skype	1.75%	BitTorrent	1.73%	Facebook	2.01%
	69.32%		*74.33%*		*72.72%*

Sources: Sandvine Global Internet Phenomena Report, July 2016

In particular, the online video phenomenon has already had a profound impact on Internet traffic. As ACG Research notes, the emergence of video streaming services has caused just about 5 percent of traditional TV viewing to shift from traditional broadcasting to streaming services over the last few years. Yet this relatively small shift in video consumption patterns has already resulted in video accounting for about 60 percent of North American fixed broadband traffic, according to ACG.

As this shift to online video advances, consumers are more and more turning to Internet-connected devices to watch the programing they want. So, instead of still relying on traditional TV sets and pay-TV set-top boxes to satisfy their video appetites, viewers are increasingly resorting to mobile phones, tablets, laptops and other portable devices to watch their content on the go, even in their own homes. Or they are taking advantage of smart TVs and game consoles to tap into the video content they desire.

Consequently, the number of Internet-connected TV devices in the home now matches or even exceeds the number of conventional pay-TV set-top boxes in the home, at least in the U.S. In a study published by Leichtman Research Group in April 2016, for example, nearly two thirds of all U.S. TV homes (65 percent) were found to have at least one TV connected to the Web in some way, up from 44 percent in 2013 and 24 percent in 2010. Due to this surge, Leichtman said, U.S. TV households now have a mean average of 2.1 connected TV devices, compared to a mean average of 1.8 pay-TV set-tops.

In particular, streaming media players are steadily streaming into homes. These small set-top boxes or dongles from such major players as Apple, Roku and Amazon are making great strides in popularity as more consumers supplement, cut or shave their pay-TV subscriptions to watch more online video. In the U.S., for example, a Parks Associates report recently estimated that nearly 20 percent of broadband households now have at least one streaming media player. Further, Parks projects that streaming media devices will reach 86 million units in annual sales globally by 2019.

At the same time, such major Internet video powers as Netflix, YouTube, Hulu and Amazon have steadily increased their consumer reach and household penetration. In the U.S., for instance, at least 52 percent of all broadband homes now subscribe to Netflix's market-leading SVOD service, with its penetration rate climbing steadily higher, according to Parks Associates. Plus, as many as 52 percent of U.S. consumers say they use Netflix at least once a week, while 27 percent use it daily, according to a recent Ericsson ConsumerLab study.

Younger consumers – those under the age of 34 – are especially flocking to OTT video, whether supplied by service providers or content providers, cutting or shaving their existing pay-TV services or, even more commonly, never signing up for conventional pay-TV offerings to begin with. In a study conducted last year by Clearleap (now part of IBM), more than 70 percent of U.S. millennials reported using a streaming service, while just 64 percent said that they have a pay-TV subscription. Younger viewers are also much more likely to use their mobile devices to view video content than older viewers.

Further, mobile video is clearly on the upswing. In its latest Mobility Report in June 2016, Ericsson predicts that mobile video traffic will soar about 55 percent annually through 2021 and account for more than two thirds of all mobile data traffic by then. This surge in video viewing over mobile devices is not limited to short-form content either. While almost 70 percent of North American mobile users say they watch short video clips at least on a weekly basis, the study reveals that almost half of them stream full-length movies at least weekly as well.

All these trends show no sign of abating any time soon. Indeed, in its latest VNI report in May 2016, Cisco predicted that IP video traffic would account for an astounding 79 percent of all Internet traffic globally by 2020, up from an already impressive 63 percent in 2015. Moreover, the VNI report projects that the world's population will watch a stunning 3 trillion minutes of Internet video each month by the end of the decade – the equivalent of 5 million years of video content per month.

As daunting as the network traffic demands are in the major markets, they may be even worse in the smaller metro markets. In a white paper last year, ACG Research noted that Tier 2 markets will handle the same traffic in five years that Tier 1 markets handle today. In other words, the Pittsburgh and Columbus of 2021 will carry the same traffic as today's Washington, D.C., and Boston, respectively.

Naturally, such higher traffic loads mean higher network transport costs for CSPs. As ACG notes in its white paper, the cost pressure is especially severe in the Tier 2 markets mentioned above because they tend to be far from regional peering points, unlike the big metro areas.

Clearly, then, service provider networks need some serious help to meet the surging Web traffic demands. In a later section, we suggest steps that providers can take to handle the traffic better, ease the strain on their networks and improve the user experience for their customers.

The Net Neutrality Conundrum

Besides the swiftly mounting traffic burdens on their networks, cable operators and other CSPs in the U.S. are facing some fresh, serious challenges on the federal regulatory front. Due to the strong net neutrality rules that the FCC adopted early last year at President Obama's behest, providers are now more limited in their ability to use some of their traditional network management tools and procedures to ease the traffic loads on their architectures.

The new "Open Internet" rules, which the federal courts have so far upheld against lawsuits brought by the cable, telco and wireless industries, spell out many of the things that CSPs can do, and especially cannot do, to manage the rising data traffic on their networks. Thus, it is incumbent upon providers to have a full understanding of the rules and their potential implications.

For one thing, the net neutrality order adopted by the FCC prohibits CSPs from blocking access to any legal content, applications, services or non-harmful devices. Further, the rules ban throttling the delivery of any lawful Internet traffic, based on the content, applications or services that are delivered. In addition, they forbid "paid prioritization" of any content, thereby preventing providers from favoring some traffic on their networks over other traffic, or offering "fast lanes" over the Internet.

These rules of the road are all clear, straightforward and objective enough. But going beyond those basic restrictions, the Open Internet rules then get more subjective and open to regulatory interpretation. That is where CSPs could easily find themselves running into trouble.

Consider the key section on "reasonable network management" practices. Although the FCC's net neutrality order generally states that all network traffic is created equal and should be treated as such, the reasonable network management section acknowledges that this may not always be possible because certain types of

services (such as video) need better treatment and/or more capacity to operate properly, especially when networks are filling to capacity. Also, the FCC recognizes in this section that different types of access networks (DSL, cable, fiber, wireless, etc.) have different technical requirements because of their different architectures.

Thus, the rules still permit CSPs to play the role of network traffic cop when they have a "legitimate and demonstrable" technical need to prioritize the data flow. But they are not permitted to do so to promote their own commercial efforts. Nor can they play traffic cop on a broad or disproportionate basis. All network management moves must be narrowly tailored to the provider's stated technical goal and must be publicly transparent as well as subject to audit. Plus, of course, all bandwidth management efforts must stand up to the scrutiny of the FCC, which is closely watching what providers do.

Or take another key section in the Open Internet order about a future conduct standard. Setting a cardinal rule for the future, that section prohibits CSPs from adopting any management practices that would "harm consumers or edge providers," such as major OTT providers like Netflix, Hulu, Google and Amazon. In particular, this section would prohibit providers from unreasonably interfering with or unreasonably disadvantaging the abilities of the edge providers to make their content, services, applications or devices available to consumers.

As these examples show, then, CSPs must tread very carefully if they wish to handle their data traffic loads more proactively by rerouting traffic around, prioritizing some content or customers over others or using other popular management practices and techniques from the past. Caution is clearly the watchword going forward.

Fortunately, CSPs do have other options available for easing the burden on their strained networks. In the next section, we spell out a promising, relatively new option that providers can leverage.

Going to the Edge: Localizing Traffic Management & Content Offloading

Given the growing Web traffic bottlenecks, the mounting network transport costs and the more stringent federal regulatory environment, U.S. CSPs could use a new service delivery model. Localized EDCs appear to offer just such a paradigm by enabling CSPs, edge providers and other content providers to bring their content, services and apps much closer to their customers, generating myriad benefits for both providers and their subscribers in those regions. The EDCs achieve this feat by enabling providers to engage in what is known as local content offloading, which we will explain below.

But before delving into the reasons why, let's start by explaining exactly what localized EDCs are. Unlike the massive, centralized data centers located in a dozen major metro areas in the U.S., local EDCs are typically somewhat smaller, neutral facilities that are now popping up in smaller Tier 2 or even Tier 3 markets to cater to a number of providers. As such, these local data centers can enable all of the participating providers to reach their customers in that market in a quicker and more affordable way.

How so? The EDCs accomplish this feat by allowing large CSPs, newer edge providers and legacy content providers to offload, or distribute, much of their bandwidth-

intensive content, services and applications to the local ISPs in the market. As a result, the most popular content, services and apps end up stored and cached at the "edger" ISP, thereby reducing the amount of long-haul traffic carried over the transport networks and cutting the data load on the Web's long-haul highways. Just as notably, the local content offloading process allows the local ISPs to deliver the desired content directly to their subscribers with no more than one router hop.

This process, in turn, produces several major benefits. First, service and content providers enjoy a significant reduction in their overall backbone network transport costs by cutting down on their long-haul carriage. In a study conducted in 2015 for EdgeConneX, ACG Research found that for a metro area with 1 million subscribers, the introduction of local content proximity could generate $110 million in backbone transport cost savings over a five-year span. That would amount to a 50 percent cost reduction over the current long-haul transport method. Thus, a modest capital investment by the CSP can save a large sum of money on transport costs.

Industry-wide, this translates into huge backbone savings and cost avoidance. In fact, bringing a majority of content to the edge for 50 million households would save the industry $5.5 billion over the course of five years. That would once again amount to a 50 percent savings over long-haul transport.

Figure 3: Cost of Backbone Transport With/Without Local Content Proximity (50 Million Subscribers)

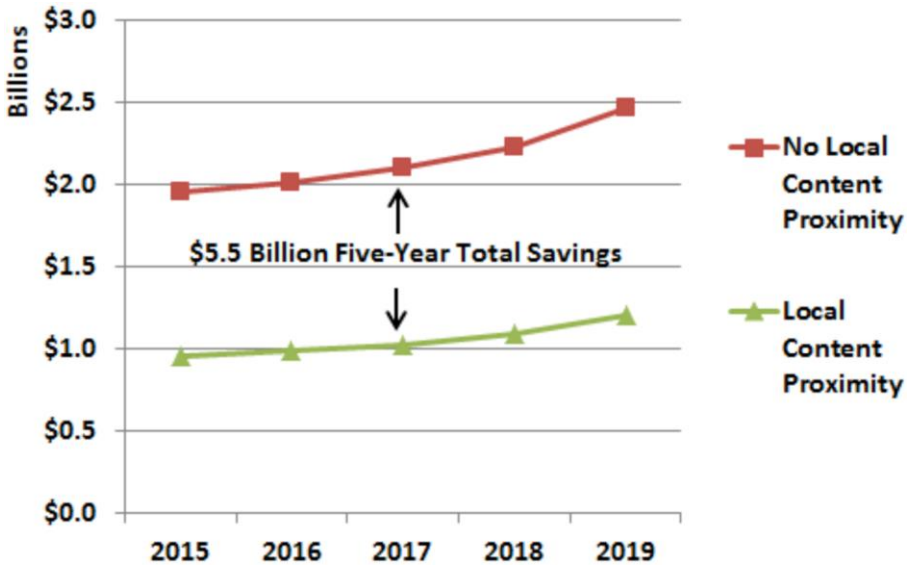

Source: ACG Research

Besides cutting transport costs, the use of EDCs can also enable providers to improve the delivery of content, services and apps to their customers in those local regions. In another 2015 study for EdgeConneX, Conviva found that two pieces of popular content experienced markedly lower service interruptions when delivered to customers locally out of the San Diego EDC: One video saw an 11 percent decrease in service interruptions, while the second saw a 17 percent decline. The videos were from two different content publishers using the same content delivery network. A competing ISP delivering the same videos from the same publishers,

but without local caching, experienced slightly more service interruptions over the same period.

Among other things, these comparative results show that standing still will cause providers to lose valuable ground in bringing a quality service to their customers. With Internet usage growing every month, doing nothing actually amounts to the same thing as doing something actively negative.

Not surprisingly, then, the end users should benefit from the use of EDCs, as well. As Conviva relates in its report, just a 1 percent increase in service interruptions leads to a 14-minute reduction in viewer engagement per video session. And when engagement levels go down, the risk of those viewers abandoning the service altogether clearly goes up, as shown in **Figure 4**.

Figure 4: Service Delivery Issues & Abandonment

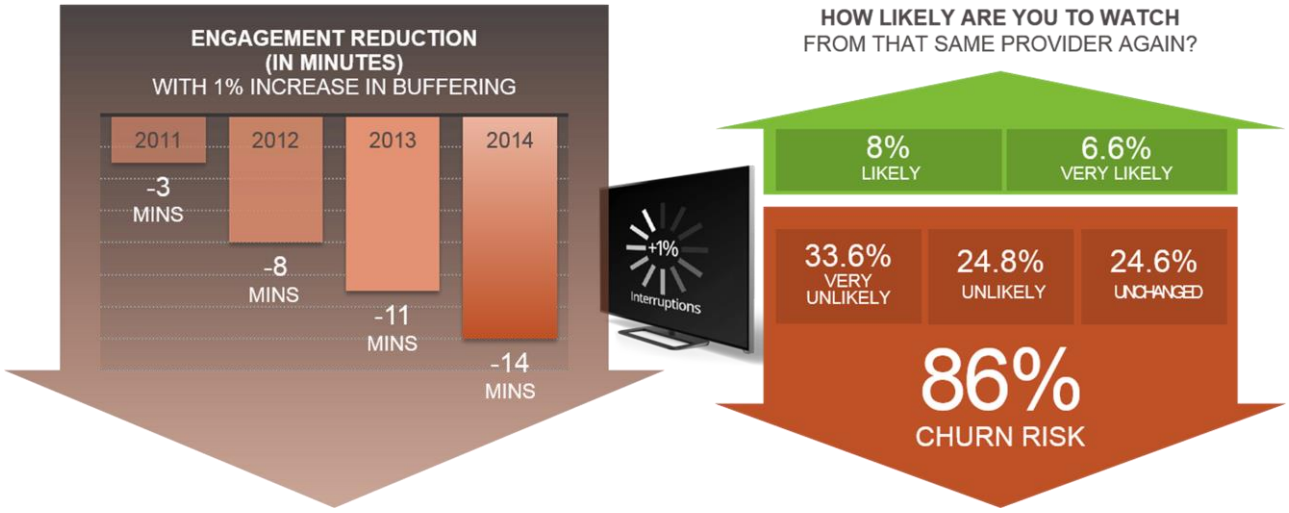

Source: Conviva, "2015 Consumer Survey Report," 300 responses

So, a decrease in service interruptions should lead to an increase in viewer engagement per session. In turn, that will lead to higher customer satisfaction levels, making them less likely to churn out and abandon their current provider. It should also lead to increased advertising revenues because of the greater engagement and lower churn results.

Similarly, in yet another 2015 study conducted for EdgeConneX, SamKnows found that the use of an EDC led to faster webpage load times for a major search engine after it moved content to a local cache in Norfolk, VA. Specifically, the average load time for the search engine's webpages dropped from 961 milliseconds to 813 milliseconds, a 15.4 percent improvement.

Furthermore, in comparative trials run in Boston, Norfolk, Las Vegas and San Diego over the past year, Conviva found that the use of EDCs by ISPs and CDNs produced lower video buffering rates for their subscribers. In Norfolk, for instance, Conviva reported that the buffering rate improved 37 percent for end users on one ISP's network, while the buffering rate improved a nearly identical 35 percent for Akamai users on the second ISP's network, as shown in **Figure 5**.

Figure 5: How Localization Cuts Buffering

	Same MSO, Two Different CDNs (One Local, One Not)	Two Different MSOs Using the Same CDN (One Connecting to a Local Cache, One Not)
Boston	Localization improves rebuffering 32%	Localization improves rebuffering 12%
Norfolk	Localization improves rebuffering 37%	Localization improves rebuffering 35%

Source: Conviva

Lower buffering rates and better scores on other customer experience measurements are important because subscribers increasingly expect streaming online video to meet the same high viewing standards as TV-based video. As **Figure 6** indicates, Conviva has found that users see QoE as the single most important factor in choosing an OTT video service – even more critical than the amount of content available or the price of the service.

Figure 6: Customers' Rising QoE Expectations

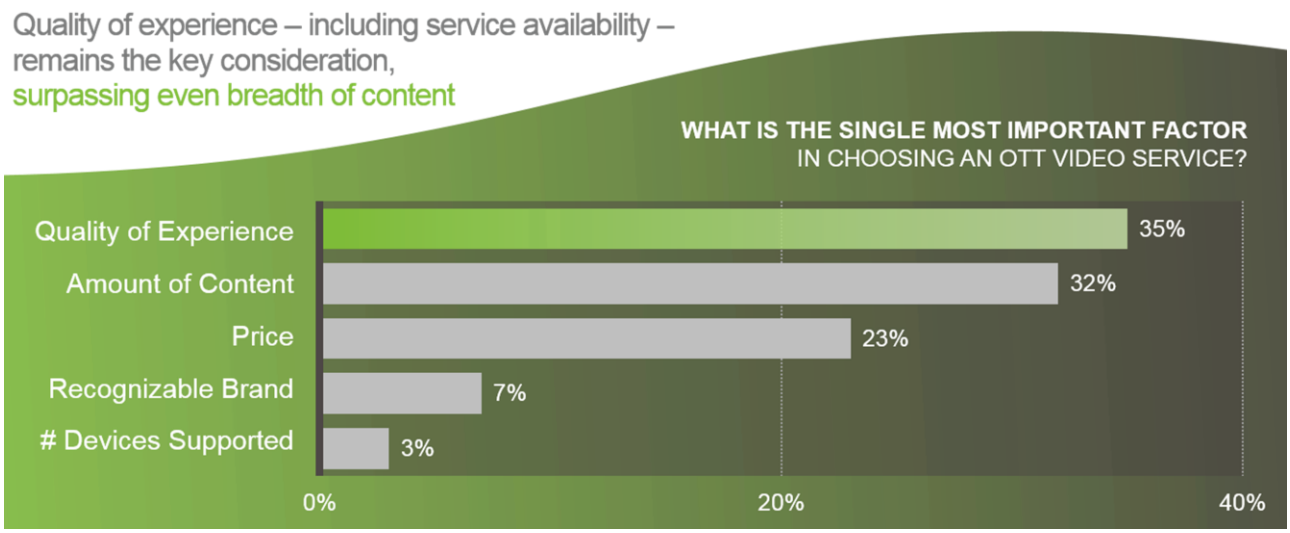

Source: Conviva, "2015 Consumer Quality Report"

Besides all these advantages, the EDC approach offers several possible benefits for business applications, as ACG Research noted in its aforementioned report. Such benefits include reducing the scale and exposure of DDoS attacks and the time needed to mitigate them, because local data centers cover fewer homes and are less complex than big regional data centers; improving the performance of cloud-based business services by reducing latency and backbone transport costs; and boosting the performance of IoT services, devices and applications by bringing the IoT devices closer to their upstream databases.

As we have seen, then, the localization of traffic management through the use of EDCs is a gift that can keep on giving. Service, edge and other content providers can reduce the peak traffic load on their heavily strained networks, cut their soaring backbone transport costs, accelerate the delivery of content, services and applications to their residential and commercial customers and improve the quality of the experience for those same customers. In turn, these improvements will lead to greater subscriber engagement, higher customer satisfaction levels and lower subscriber churn rates, even in an increasingly competitive environment.

Conclusion

As we have seen, traffic congestion is growing exponentially on both the public Internet and private CSP networks. Due to a potent combination of more Internet users, increasing number of Web-connected devices, ever-faster broadband speeds and the growing shift to IP video, content and service providers alike are increasingly struggling with clogged networks, less capacity, network slowdowns, more service interruptions and other quality-of-delivery issues. Plus, long-haul network transport costs are continually mounting, forcing providers to shell out ever-larger sums on network upgrades and new equipment.

With the U.S. government's recent adoption of strict net neutrality rules, the problem could only get worse for providers of all stripes. Stripped of the ability to use at least some of their traditional data management tools and techniques to control, prioritize and direct the traffic flow across their networks, they are now more limited in what they can do to ease the growing bottlenecks. Yet they are still the ones held primarily responsible by customers for the QoE problems that the bottlenecks produce.

Fortunately, there is a valuable new alternative increasingly available for cable operators, telcos and other CSPs. As spelled out in this paper, that alternative is the use of localized EDCs. With this approach, providers can offload their content in data centers much closer to their actual subscribers, caching it in the local data centers until it's actually ready to be viewed. That way, the content can be delivered quickly, easily and more smoothly to customers without running into traffic bottlenecks and the QoS and QoE problems that they invariably cause.

By adopting this strategy, CSPs can also slash the growing transport costs on their backbone networks. As estimated by ACG Research, those backbone savings could amount to as much as $5.5 billion over five years for the entire U.S. cable industry, generating up to 50 percent in overall cost savings.

In other words, going local and getting close to the edge really can make the difference. The sooner providers can start doing that, the better off they will likely be, along with both their residential and commercial subscribers.