Empowering Internet Users to Manage Broadband Consumption

Peter Sevcik
June 2012

It's time to empower Internet users with bandwidth consumption transparency. Users have the right to know how websites and applications are consuming their bandwidth--especially since it is users who are footing the bill for the bandwidth. Content providers, advertisers, and vendors owe it to users to be prudent in the amount of traffic they generate, and to give users at least some ability to determine what consumes their bandwidth. This paper proposes that standardized bandwidth miles-per-gallon-like consumption data be provided for devices, services, and applications, empowering users to make informed buying decisions.

NetForecast Report 5109

What is Bandwidth Consumption?

There are many ways to quantify Internet traffic. Some methods count users, while others count cookies, registered accounts, ads, locations, clicks, page views, eyeballs, conversions, or revenue. Each of these methods counts different things as "traffic". Underlying all of these concepts of traffic, however, is bandwidth. Bandwidth is the fuel that powers the Internet by transmitting bits from point A to point B much as gasoline is the fuel that powers a car to move from A to B.

ISPs provide users with the bandwidth “fuel” and Internet-enabled functions consume that bandwidth. By connecting users to the vast Internet, ISPs provide an audience and economic foundation for the commercial Internet's existence.

In the early days of the Internet, when users got online via dial-up modems, applications and services were designed to work efficiently to consume limited bandwidth. Overall bandwidth consumption was also limited since each time someone in the home dialed into the Internet they did so with just one PC running one application. Applications used less bandwidth and the contention for the bandwidth that was available was limited.

Broadband to the home changed that, giving the impression that the user’s Internet connection has unlimited access capacity. Many new applications, uses, and businesses arrived, each of them seemingly free to paint on an almost limitless canvas of bandwidth. Many of these applications and services operate with the notion that they are free, running on unlimited bandwidth—a perceived utopia of costless transactions. But in reality, Internet access is never costless or limitless.

How Applications Consume Capacity

Many people think that when they purchase a certain speed of Internet service, that speed is their capacity to consume, and as long as any use of the speed is less than the purchased speed then all will be well. This is not the case because capacity is not speed.

Capacity is bandwidth purchased over time, as shown by the outer box in Figure 1 below. Time may be a day, a month, or a billing period. The capacity is the total area of the outer box typically stated in MB or GB per month. In Figure 1, several users and applications that are consuming available capacity are depicted by colored boxes. Consumption in this example is the sum of the area of all colored boxes relative to the capacity that the user purchased.

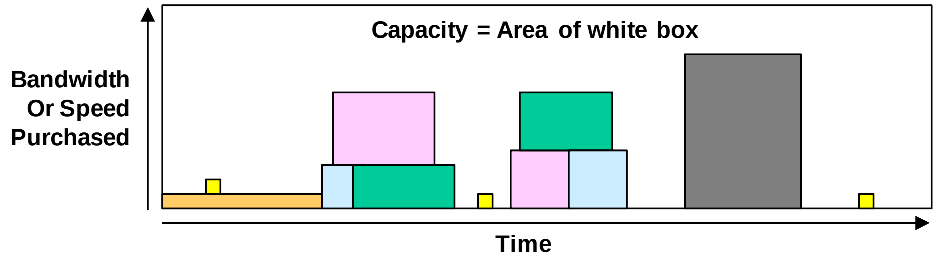

Figure 1 – Applications Consuming Bandwidth Capacity

Each application consumes a slightly different amount of capacity, and when multiple applications operate simultaneously, their use of "speed" is summed so boxes stack on each other. The large gray box application utilized a lot of both bandwidth and time. The sum of all the applications consumes 30% of the capacity in Figure 1.

All Capacity is Fixed

Every Internet subscriber has purchased a fixed amount of Internet capacity. The capacity of unlimited plans is the product of the bandwidth purchased and time. For example an unlimited 1 Mbps service has a total monthly downstream capacity of just over 300 GB if used at maximum bandwidth at all times. Some ISPs institute a usage cap which limits capacity to a fixed value.

Regardless of how an ISP service is delivered or sold, each user's purchased *capacity is always a fixed resource*. The remainder of this report discusses how that fixed capacity is consumed.

Understanding Traffic Categories is Key to Managing Consumption

Given the explosion of Internet content, services, applications, and connected devices users need to understand how all of these interact to consume the bandwidth capacity they have purchased. To do so, however, one must first understand how these things are currently working in the broadband Internet.

This paper explains how producers of Internet content, services, and applications are driving bandwidth consumption in ways that are not visible or controllable by users, and it proposes to start a dialogue to improve transparency. Because bandwidth is neither limitless nor costless, users must be empowered to manage their broadband consumption.

There are three primary traffic categories that drive bandwidth consumption on the Internet today: content, advertising, and background traffic.

Content

After a user clicks on a button or enters text they expect to see new content within a reasonable elapsed time (i.e., page load time). The content can generate a wide range of direct consumption such as:

- A few bytes of capacity when updating a stock price within a second, or
- A few gigabytes of capacity required to watch a movie over 2 hours.

The content is delivered by an enterprise operating over the Internet that has its own business model for giving users access to the content.

Advertising

> "If you are not paying for it, you're not the customer; you're the product being sold."
>
> - Andrew Lewis
> August 26, 2010

The predominant content delivery business model is "free." A free service is supplied by a company that gives away its service to users while being paid by advertisers. The companies make money by selling access to their users. Their customer is the advertiser and the user is the "audience." The objective of this business model is to generate clicks, which deliver more ads that require more bandwidth for which the audience (user) pays.

Since the free advertising-based model is the most prevalent in the Internet commerce space, it is important to understand its technical and economic implications. Radio and TV networks were built upon advertising, with free service to the audience. The volume of advertising a broadcaster sent to the audience did not affect the performance or the cost of the user's TV. It relied on FCC-regulated broadcast spectrum which was free.

Today, advertising load over the broadcast spectrum ranges from low for public broadcasting, to moderate for typical commercial television, to high for infomercials. It does not matter which of these broadcast types a consumer watches – there is no additional cost to the consumer and no impact on their TV's performance. Generally newspaper and magazine readers experience the same flat fee regardless of the number of ads in the publication. Other than the more expensive Sunday newspaper, the thick issue does not cost more than the thin issue containing fewer ads.

But the same is not true for the Internet-based, advertising-supported business model. Since bandwidth is Internet fuel for which the user pays on a subscription basis to their ISP, the free, advertiser-supported model isn't free at all – it costs the user in the form of increased hidden bandwidth consumption for which the consumer pays their ISP.

Background Traffic

The third, and fastest growing, category of traffic is traffic that the user never sees. This traffic is generated by behind-the-scenes services used by devices or applications to communicate information to third-parties which has not been directly requested by the user. The communication can be used to improve or safeguard a device, service, or application. This secret, and almost always hidden (often encrypted), traffic can be used to update software and modify designs for hardware vendors. In some cases it's used to mitigate security vulnerabilities constantly arising in software or devices. The most widely known example is periodic software updates.

In many cases background traffic is used to track product usage and sell that usage information to third parties. This type of background traffic forms the underpinnings of the advertiser-supported business model.

Figure 2 shows a real-world example of background traffic. It is the daily traffic consumed at my house over a month when no one was home. The Internet connection remained on with only one PC connected to the network. The PC was configured not to receive any software updates or reach out onto the network in any way. Nevertheless, an average of 70 MB per day traversed the broadband connection – and no one was home! This shows that the background traffic is essentially impossible to stop as long as the hardware is connected to the Internet.

A newer class of background traffic has recently emerged that pushes content to users and pre-loads it to devices in case the user wants to see it. However, in many cases the user never sees the vast majority of this traffic, despite the fact that it is being transferred over the user's broadband access line to the device or browser. As iPads, iPhones, Nooks, Kindles, and other mobile devices multiply and join the established population of desktop and laptop computers, the nature of traffic to and from all device types is changing. The key reason for this change is that users are downloading apps, which replace the browser for completing most tasks. Mobile device users generally rely more heavily on apps to do their bidding, while computer users continue to rely more on browsers. To the surprise of many, app-centric activity consumes more bandwidth than browser-centric activity.

Given that mobile device users are more likely to pay for bandwidth consumed, understanding how apps eat bandwidth can avoid end-of-month billing surprises. However, many mobile devices also consume wireline broadband services when they switch to Wi-Fi access in the home.

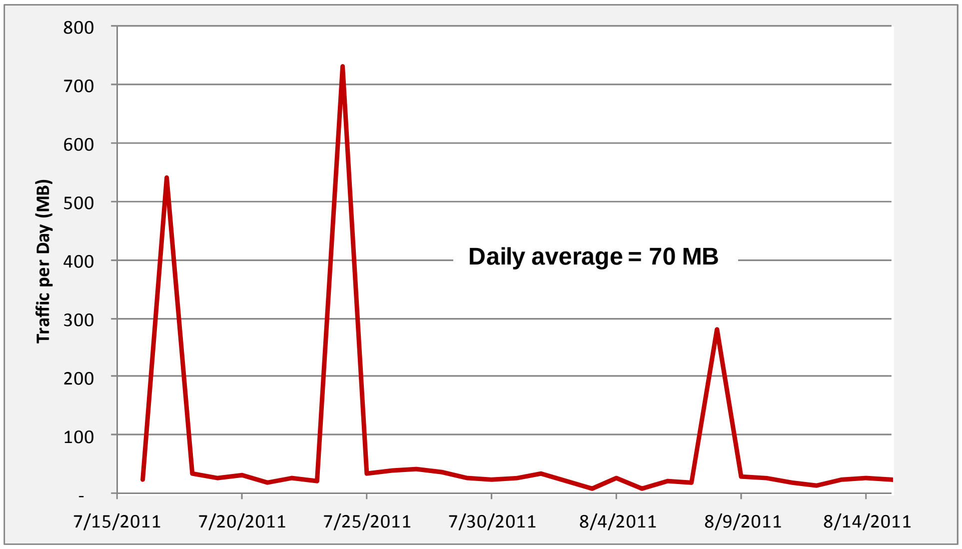

Figure 2 – A Snapshot of Background Traffic

Browsers generally deliver real-time content that is immediately visible on the screen (e.g., one WSJ article). Apps, on the other hand generally download a lot of content to the user's device that a user may or may not ever view (e.g., 124 WSJ articles in a typical day's issue). For example, comparing two basic ways to read the Wall Street Journal:

- One article as a web page on Safari browser on iPad: 2 MB
- The WSJ application daily issue on iPad: 58 MB

The application is programmed to deliver the entire paper, regardless of whether the user intends to read every article every day. Because the cost of transmission is zero to the application creator, the application operates without regard to the impact on the network – and without regard to user preferences. Another way of understanding the difference between the two delivery methods is to consider that in a day the user will have to read 29 Wall Street Journal articles for the traffic generated by the app and the browser to be equivalent. If the user reads fewer than 29 articles, then the browser method was more efficient.

If the Wall Street Journal loads 124 articles each day, and if a user views 6 of them, then only 5% of the consumption (traffic on the user's access line) is classified as content. The remaining 95% is background. An issue that downloads to a tablet but is never viewed represents 100% background traffic.

The WSJ app is a more convenient experience than reading the paper via a browser. After all, apps are a popular way to read a large variety of newspapers and magazines. The user is paying for that convenience by making a choice to consume more of their purchased capacity. They just need to understand the choice they are making.

The Three Traffic Categories on the Web

Figure 3 shows the traffic categories seen on web pages since 1995. We use the Keynote Business 40 as the proxy for a typical business web page. NetForecast has been profiling the then-current 40 web pages since 1995. We also use data from the HTTP Archive to help determine the mix of traffic categories.

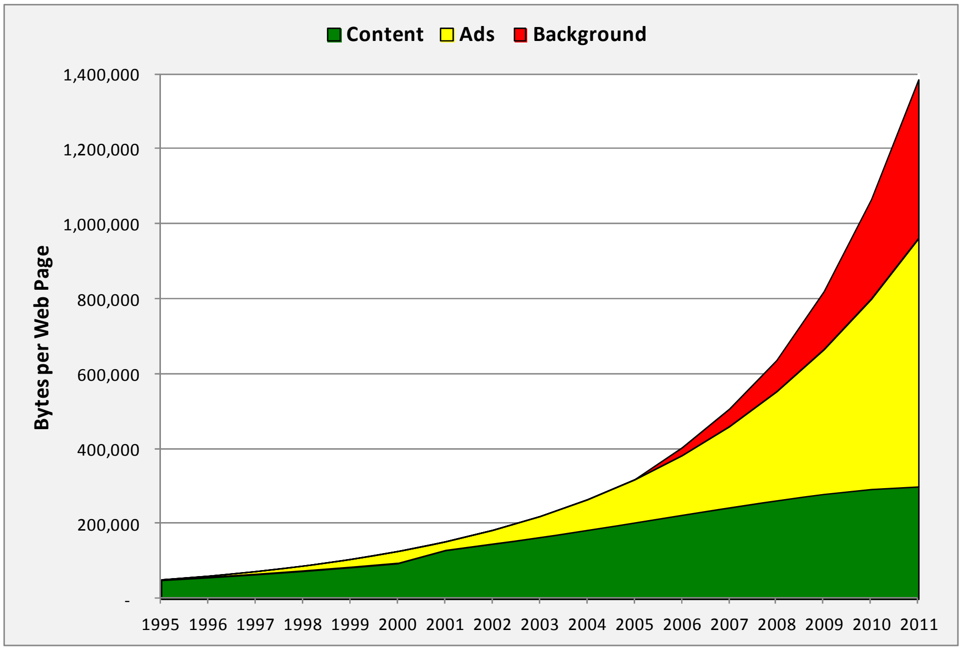

Figure 3 – Traffic Category Mix of the Web Over Time

The graph above shows some interesting patterns. The rate of content growth has slowed dramatically, while the other traffic categories continue growing at a rapid pace. Current studies show that a typical web page needs about 1 MB to load (sum of green and yellow). But those who track web traffic have not taken the background traffic directly associated with web pages into account. Background traffic is a relatively new phenomenon driven by cookie processing and browser pre-fetching. It started in 2005 with the introduction of Google's Chrome browser.

Based on our analysis the current mix of web traffic is approximately:

- Content 20%
- Advertising 50%
- Background 30%

The typical web user looks at or interacts with only one-fifth of the bytes passed over their Internet connection. The other 80 percent is either advertising or unsolicited background traffic. The user, of course, has to pay for the entire 100% of these transmissions. A more alarming statistic is the compound annual growth for each category as of the end of 2011:

- Content 2%
- Advertising 30%
- Background 60%

Background traffic, about which the user is in the dark, is growing at an alarming rate. Soon users will be paying to move nine bytes for every one they are using.

Application Efficiency

Many services on the Internet can be supplied within a consumption range. NetForecast calls this application efficiency since it supplies essentially the same service with less consumption.

Netflix did this in March 2011 by offering a lower bit-rate to supply movies and TV shows to Canadian subscribers. Neil Hunt, Netflix Chief Product Officer, announced that Netflix was suddenly using 2/3 less data on average, "with minimal impact to video quality." They introduced a new feature within the "your account" settings that permits subscribers to choose among the following video quality settings:

Good quality (up to 0.3 GB per hour)
Better quality (up to 0.7 GB per hour)
Best quality (up to 1 GB per hour or up to 2.3 GB per hour for HD)

Netflix offered this feature in March 2011 because some of the Canadian wireline broadband ISPs have very low usage caps. At the time the lowest service tier usage cap at Rogers Communications, was 2GB. Since then, the lowest cap is 15 GB per month for their least expensive service.

What Netflix did was to set the default service at the "Good" level for all Canadian subscribers. The user can change it to the more aggressive settings if they choose. This feature is now also on all US Netflix accounts but the default is set at Best Quality. In the US, the user must *choose* to spend less on bandwidth.

Splitting Up a Fixed Pie

A lot of advertising and background traffic consumes the fixed broadband pie in every home. Figure 4 below shows the relative percentage of such traffic consumed by five general activities. The relative ratios of content, advertising, and background shown are derived from our measurements and the use case descriptions in this report.

Of course, there are many different applications, websites, and devices that consume the three traffic categories with very different ratios. We do not presume to know the exact ratio of any specific Internet device or service. The fact that this is difficult to discover is the point of this report. We have however, observed that the ratios for each device or service is stable for a reasonable period of time. The ratios are a byproduct of how the device or service is architected, characteristic which changes infrequently.

We offer the following figure as a general snapshot of the ratios in 2011 based on our measurements. If your website, service, or device has a different ratio, then please publish it somewhere. Better yet, put it into a label proposed at the end of this report.

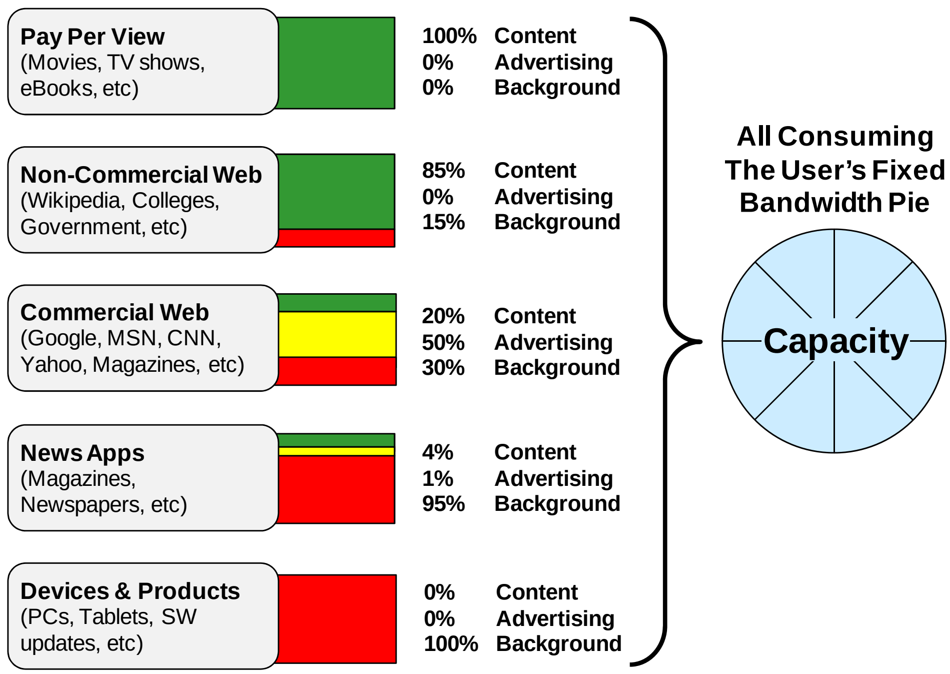

Figure 4 – Examples of Demands on Users' Fixed Bandwidth Capacity

The five activities are ranked by the relative percentage of user-visible content (shown in green). Estimating total use by activity is difficult because little hard data exists and many activities vary by how often or how much users participate in the activity over a month, thus the total consumption of each activity is not shown.

All Internet activities rely upon consuming a portion of an Internet subscriber's fixed Internet bandwidth allotment. This is a fixed-size pie that is split many ways. Under current business models, users essentially subsidize the advertising and background traffic (shown in yellow and red) that they receive without asking for it.

Notice that devices and apps, the newest rage on the Internet, have the least percentage of green. Devices like the iPad and Android tablets, along with their apps, are essentially cost-shifting their upkeep to consumers.

What the User Needs

Total bandwidth and individual subscriber speeds will grow, but at any given time a subscriber has only a fixed amount of purchased Internet capacity available. The subscriber purchases capacity as a resource from the ISP. Many Internet-based businesses assume that the user has purchased a lot of (or unlimited) capacity, and that the business is entitled to consume as much of that capacity as it wants.

At the same time, subscribers are busily adding multiple users in the home, devices, and applications, often running simultaneously. The fixed-size pie of bandwidth should be understood and properly sliced among the needs as defined by the user. However, the growth of background services takes that power away from users. In essence, these hidden services are invisibly hogging the user's capacity. And, while this may have seemed harmless when broadband networks seemed limitless, it becomes a real burden as it impairs the user's experience more and more.

For example, if they purchase paid-for content such as a movie, other uses in the household or other applications in automatic transmission mode may consume so much bandwidth that the paid-for content experience is not enjoyable. Is it right that at the same time you are watching a Netflix movie that you paid for, Microsoft can impair that experience by upgrading Office on a PC in another room?

We do not propose that the business relationship between the consumer and ISP change in any way. We do propose that the consumer be supplied with the information and controls needed to make informed choices. To do this, we propose the following tools that help consumers make informed decisions about their broadband capacity resource consumption.

Demand for Internet "fuel" has another parallel with gasoline. The *quality* of gasoline is defined by octane standards determined by the government. That is analogous to ISP bandwidth as measured and publicized by the FCC via the Measuring Broadband in America Program.

Taking the gasoline analogy further, now that we have a way to measure fuel quality, we need a miles per gallon (MPG) type label on software, services, and devices that consume bandwidth.

Consider the following examples.

A consumer looking to buy an eye-catching car like the Ferrari 458 Italia shown here can go to a federal government website to learn that it delivers only 14 miles per gallon using premium gasoline. This consumption label is also on the window of the car in the showroom. These labels also inform the buyer that the purchase will include a gas guzzler tax. Clearly, we have decided that consumers need to be warned.

The Apple iPad 3 is similarly eye catching. The new Retina screen resolution exceeds the best HD televisions. This is achieved by rendering all previous single pixels with 4 finer pixels. This means that most graphic or media content requires 4 times the bandwidth capacity to deliver the same content. Where was the consumption label that should have warned consumers *before* they purchased the device? In fact, this is not just a theoretical issue. Since the launch of the iPad 3, consumers have reacted negatively to the impact it has had on the consumption of wireless broadband bandwidth. At the time of this writing, consumption information is just now disseminating via the press after Apple has sold more than 3 million iPad 3s.

Consumers need information when they are making a decision that will immediately and seriously impact their bandwidth consumption. For example, Netflix should let users know when they are ready to start a movie that they have delivery choices that will affect the total consumption the movie will require. The current video quality settings feature supplied by Netflix is a global account setting meaning that everything the user downloads will be delivered at a pre-set quality. This is good start but an improvement would be to give the user the ability to change the quality setting at the start of each movie. Such a message might read:

> *You are about to start a 2-hour movie at a resolution that will consume 2 GB of your broadband service. You have the following options that will consume less.*

The Apple App-Store already does this on the iPhone. If you are connected via a cellular data plan (Wi-Fi is off) and you decide to update several apps, the user interface says:

> *You are about to update several apps which will consume a lot of traffic on your cellular data plan. Are you sure you want to do this or do you want to turn on Wi-Fi access?*

Providing such information is being a good network citizen and helping consumers make an informed choice. All major bandwidth consumption decisions should come with such a warning, along with suggestions about ways to do the job that will use fewer resources and, in some cases, work better. For example, switching to a home Wi-Fi network for the iPhone updates saves money because the retail cost of wireline broadband access is significantly less than wireless broadband access.

An Internet Consumption Labeling Proposal

We propose the following as a prototype of the Internet Consumption Label. This example is based on current knowledge of typical Internet browsing. Of course every supplier of an app, device, software or web service will have their own label and the specific information presented will be dramatically different.

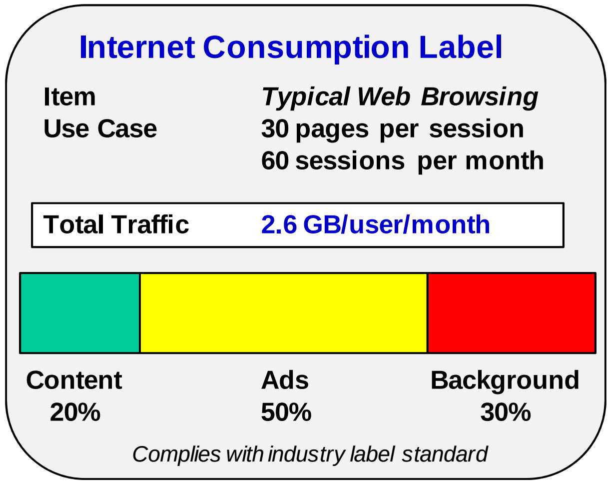

The specific data presented on the labels should be standardized by an industry group to ensure uniformity. The group should also educate consumers on how to interpret the label and make informed choices regarding their personal Internet consumption.

The industry already has the information to fill out an Internet Consumption Label. Internet usage is the most instrumented, followed, and analyzed technical activity. User tracking is very well understood. Websites and app suppliers know much more than what would appear on the label because they use that information to sell goods and services as well as advertising. The information required for the consumption label, for example, is a small fraction of the data Google gathers with Google Analytics. The change proposed is that the information be made available to consumers for their benefit. The need for transparency is urgent. The question is, who will lead the way?

About the Author

Peter Sevcik is President of NetForecast and is a leading authority on Internet traffic and performance. Peter has contributed to the design of more than 100 networks, including the Internet, and is the co-inventor of two patents on application response-time prediction and congestion management. He works extensively with the SamKnows system in support of the FCC Measuring Broadband America project, analyzing operational integrity and performing deep data analysis. He can be reached at peter@netforecast.com.

NetForecast helps enterprises understand and improve the performance of networked applications.

Additional information is available at: www.netforecast.com