NetForecast

Want to Understand Why You Hit Your Usage Cap?

Peter Sevcik and Rebecca Wetzel
May 2018

Introduction

If you want to understand—and possibly even count—the traffic crossing your household's Internet connection, there are some things you should know. Many people believe they have a firm handle on all the major traffic sources in their household, but even the most tech-savvy among us can be surprised by "mystery" traffic. This guide will help you understand unexpected traffic, and it will help do-it-yourselfers who choose to do their own counting avoid common pitfalls.

Understanding Unexpected Traffic

If you examine your Internet usage closely, chances are you will see unexpected bytes register on your meter. But if you're like most users, you don't know how much traffic your household produces. So, for starters, it is useful to understand general Internet usage trends because they are likely to provide insight into your home's usage. The most important trend is the steady increase in video content consumption over the Internet. Streaming services now constitute the lion's share of most homes' Internet usage.

How do we know?
NetForecast audits Internet Service Providers serving over **100,000,000** US broadband subscribers

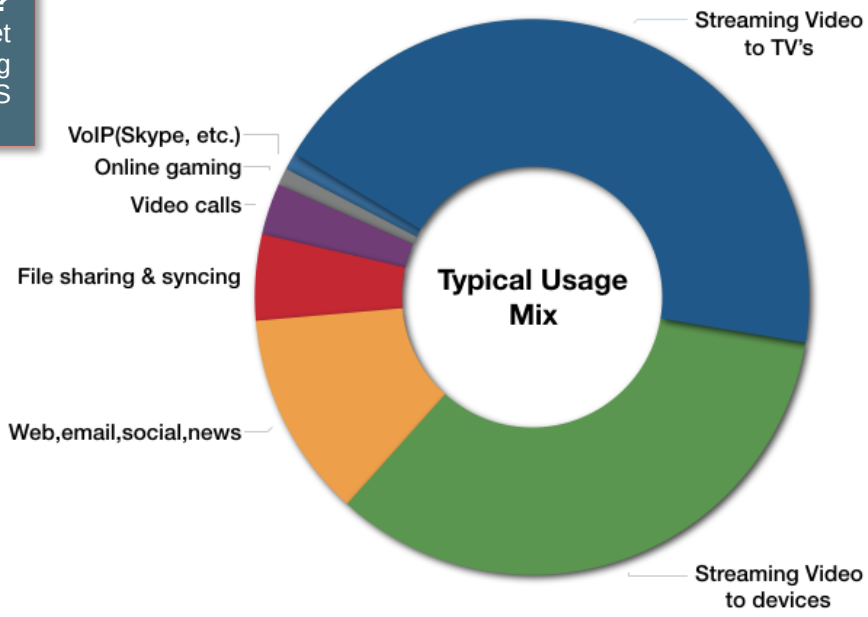

Figure 1 – Data Usage Mix in Typical US Households
(Data source: Cisco Visual Networking Index)

Today, most sites and services are delivered via streaming. A single user action often triggers a constant stream of content that lasts a long time. For example, you may check social media for 5 minutes and move on, but the app keeps streaming on your device for as long as 20 minutes.

Video streaming services do this on a grand scale. At the end of a TV episode you have watched, the service often begins the next episode within a few seconds without your "permission," thus using more data. Have you ever returned to a device later to watch the next episode in a series and the device starts two episodes later? Those additional episodes were streamed to your device while you were away. This can add hours of video traffic to the single hour you actually watched, adding reams of data to your Internet usage meter without your knowledge or approval.

Your devices on autopilot

In the past user action was needed to **start** data traffic.

Today user action is needed to **stop** data traffic

Not long ago, the Internet required human user action to initiate data usage. When you read an email, clicked on a web page, or opened a news article, your action caused an expected transfer of information. User action was followed by Internet reaction. By the end of a day users could gauge their Internet usage. That's no longer the case. In addition to video streaming, here are some other services that can slip under your radar.

Internet of Things - One common surprise is the number of traffic-generating devices in a home—from PCs, smart phones, tablets, digital video recorders and printers, to thermostats, appliances, door locks, game consoles, and cameras. Many non-PC devices "phone home" to a manufacturer or support service. For your convenience, these automated connections are invisible to the user, making it difficult to gauge the traffic generated.

File syncing - NetForecast's research has revealed unexpectedly high consumption caused by cloud storage and file-sharing services such as Google Drive, OneDrive, Box, Dropbox, and iCloud. Some cloud services can enter an endless synchronization loop, which dramatically increases traffic counts. We have observed this when very large files are shared across multiple users. If you use one of these cloud services and experience extraordinarily high data consumption, we suggest investigating your configuration with the help of your cloud service provider.

Out of control updates

NetForecast discovered a home device that constantly requested a software update during the first 7 days of alternate months, generating more than 500GB of data usage every other month.

The server for the device was not responding as the device expected and the manufacturer was unreachable. The only way to stop the usage was to disconnect the device.

Wi-Fi - Someone outside your home may use your connection without your permission or knowledge. One of the first things you should do is change your Wi-Fi password if you suspect unauthorized usage. This is common in densely populated urban areas where hundreds of devices can see your Wi-Fi signal and potentially attach to your network.

Security - Security-related issues can lead to unexpected traffic. A PC can be hijacked and generate traffic that has nothing to do with any user in your home. Specifically, botnet and malware-infected devices can be leveraged for outbound Distributed Denial of Service (DDoS) attacks against targets on the Internet, and/or can be used as proxies to route traffic for bad actors. Bad actors can also use User Datagram Protocol-based amplification attacks to exploit vulnerabilities on a home router, resulting in high data consumption [1].

Software Updates - Most popular software has automated update features that download and install updates. This automation is for your convenience and protection, but the traffic it generates may come as a surprise. Although each program update download may be small, when you multiply a modest download by the number of programs calling for updates and the number of devices in your house, such traffic can be substantial. Furthermore, in some cases vendor default settings are aggressive, checking and downloading every possible option every hour, even if they are unneeded

(e.g., a software program may automatically load its interface in a dozen languages for a monolingual household).

We recommend that you check your software settings and align update size and frequency to your needs, bearing in mind the amount of traffic generated.

Additional information about hidden traffic is available in the NetForecast report *Empowering Internet Users to Manage Broadband Consumption* [2].

Useful Information If You Want to Do Your Own Counting

If you wish to perform your own Internet Usage meter validation testing, it is important to understand factors that may cause your measurements to vary from the ISP meter results.

Avoiding binary versus decimal math confusion

ISP data usage meters report in gigabyte increments, so if you are measuring your own usage, make sure you are using binary math. One gigabyte is a binary number not to be confused with one billion bytes. The following table illustrates the danger of applying decimal notation to byte counts.

Binary					Decimal	
KB	Kilobyte	1,024	≠	Thousand	1,000
MB	Megabyte	1,048,576	≠	Million	1,000,000
10 MB	Megabyte	10,485,760	≠	10 Million	10,000,000
100 MB	Megabyte	104,857,600	≠	100 Million	100,000,000
1000 MB	Megabyte	1,048,576,000	≠	1000 Million	1,000,000,000
1 → GB	Gigabyte	1,073,741,824	≠	Billion	1,000,000,000
2 → TB	Terabyte	1,099,511,627,776	≠	Trillion	1,000,000,000,000

Here are some typical errors introduced by binary/decimal confusion:

Note 1: One GB is 2.4% larger than 1000 MB. Many people mistakenly believe that 1000 MB is the same as 1 GB. It is not. The reason for the confusion is the mixing of binary and decimal math.

Note 2: One GB is 7.4% larger than 1 billion (pink vs. blue in the table above).

Where you measure matters

You can gather your own usage information either from a computer or from the network on your premises. A computer can track what is downloaded to/uploaded from it, but it does not report network protocol overhead because such data is hidden within the PC operating system (you need special instrumentation software to see all the protocol traffic).

For example, if one looks only at the size of a file on a PC (payload), that value does not include any protocol overhead that is required to transmit/receive that file, which may lead one wrongly to conclude that the ISP meter is overcounting.

If you measure traffic at the network layer, you will see the payload traffic plus overhead from protocols like TCP/IP and Ethernet, which generally add about 6% to 9% overhead to the payload traffic for large packets and a larger percentage for small packet traffic like VoIP. The meter system counts the traffic as seen on the wire, which includes the payload plus protocol overhead, so it should closely match the network view. As an example, a simplified data packet would look like this:

Some other notes on self-counting

Network layer counting is best done using an intelligent switch or router.

Be aware however, that these devices often fail to count all protocols (e.g., Ethernet), so you may be undercounting. It is important that your network device counts ALL traffic passing into and out of the Internet, and that your device does not count local traffic (e.g., traffic to printers or local music servers). You must be certain to count all Wi-Fi traffic to/from the Internet. You must be careful to configure your measurement software to count only the relevant traffic.

Doing your own counting also requires careful data gathering.

Switch and router counters typically default to zero when the device boots, and subsequently display cumulative usage counts. These counts continue to increment past ISP billing month boundaries. To track your ISP's usage meter accurately, you must record counts periodically—especially at the billing date boundary. Keep in mind that the date boundary depends on the time zone your ISP uses.

Details that may seem minor can alter your counts.

For example, we are aware of several router models that appear to count properly, but only count usage for devices in the DHCP table at start up. Usage from devices added to the network after the router booted went uncounted. Rebooting the router brought the new devices into the counts from the reboot onward. The subscriber reasonably concluded that the ISP was overcounting, but in fact, months had passed since the last router reboot, and new devices were introduced into the home during that period. These new devices generated significant usage that the router did not count but the ISP did. Properly measuring home usage requires technical know-how, careful attention to process, and patience.

Conclusions

Unexpected traffic sources can cause your household's Internet usage to jump. If you are subject to a data cap, it can be useful to understand and manage your usage to remain below the cap. In addition to keeping an eye on obvious traffic sources such as video consumption, it is important to pay attention to such things as file syncing, Internet-connected "gizmos," software updates, unauthorized network use, and malicious activity.

If you are interested in taking your own usage measurements, make sure you educate yourself about the ins and outs of what, where, and how you count—because without that information you may draw false conclusions.

References

1. US Computer Emergency Readiness Team (US-CERT), *UDP Amplification Attack*
2. Sevcik, *Empowering Internet Users to Manage Broadband Consumption*, NetForecast Report 5109, presented at The Future of Internet Economics, Technology Policy Institute, June 15, 2012.

About NetForecast and the Authors

NetForecast independently audits the accuracy of ISP data usage meters to ensure subscribers are billed fairly and to enable ISPs to correct system inaccuracies. During the past decade we have audited ISPs serving 73% of US broadband-connected households, and we have established industry standards for data usage meter accuracy. Our methodology has been independently audited and approved by NERA Economic Consulting.

About the Authors

Peter Sevcik is the founder and CTO 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 three patents on application response-time prediction and network congestion management. He pioneered Internet usage tracking techniques, and invented the Apdex methodology. He can be reached at peter@netforecast.com.

Rebecca Wetzel is a principal at NetForecast, and a data communications industry veteran. She helped realize the commercialization of the Internet in its early days and worked to design and market some of the Internet's first value-added services such as IP-based VPNs, web hosting, and managed firewall services, as well as Internet protocol implementation testing services. She also spent many years as an Internet industry analyst and consultant

Additional information is available at: www.netforecast.com