Accuracy Audit of Cox Communications' Internet Data Usage Meter

By Peter Sevcik
May 2015

Cox Communications has deployed a data usage meter system to measure and report how much Internet data traffic a residential subscriber consumes and generates. Cox reports Internet subscriber usage information on its customer portal, and notifies subscribers via email when they approach the monthly data usage limit associated with the service package they have purchased. Cox engaged NetForecast to perform an independent audit of the accuracy of its data usage meter system.

NetForecast generated test traffic, performed independent traffic measurements, obtained daily usage meter records from Cox for three test locations, and compared NetForecast's usage data with Cox's data usage meter records. NetForecast performed tests in three locations in September and October of 2013, and retested one location in October and November of 2014. Cox had no prior knowledge of the timing or volume of the NetForecast test traffic.

Cox established an accuracy goal for its Internet data usage meter to correctly measure traffic passing through a subscriber's cable modem within plus or minus (+/-) 1.0% for each month. Our analysis validates that the meter is accurate within plus or minus (+/-) 0.1% (one tenth of one percent) for the month, and within plus or minus (+/-) 0.2% (one fifth of one percent) for the cumulative daily sum at each location—an outcome that falls well within Cox's 1% stated goal.

The Cox Data Usage Meter

Cox's Internet data usage meter provides subscribers with information about how much traffic has crossed their residential Internet connection. Subscribers can view a data consumption summary for the current period in a "gas gauge-like" view (Figure 1). They can also view more detailed usage reports as shown in Figures 2, 3, and 4.

All of the reports sum traffic passing from the Internet into the household (downstream) with traffic passing from the household (upstream) to the Internet into a single value.

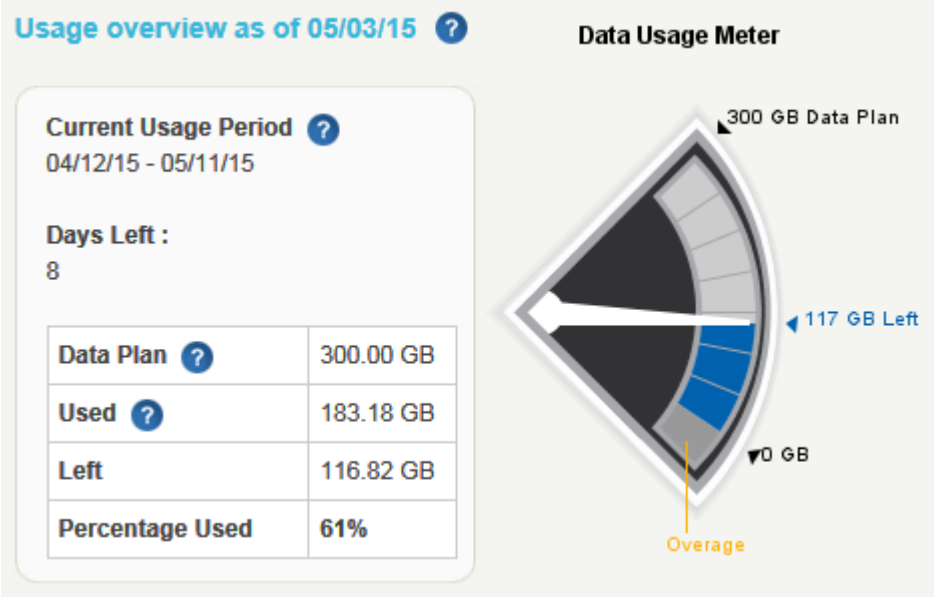

Figure 1 – Sample Subscriber Summary View of Usage

NetForecast Report NFR5118

Cox's Internet usage reports are accessible via the Cox customer portal at http://www.cox.com/. Subscribers with online accounts can access their Internet usage information by logging in at the sign-in prompt at the top left of the portal menu bar, and highlighting **My Connection**. This will open a navigation window. The **Data Usage Meter** link is located under **Internet Tools** in the bottom section of the navigation window.

By scrolling down the data usage meter page, subscribers can view a detailed daily usage graph for the current or previous usage cycles (Figure 2), and can choose other data usage views, including: a table of daily consumption values (Figure 3); a cumulative monthly graph showing how much traffic has been consumed relative to the subscriber's monthly usage limit (Figure 4), and a monthly history graph.

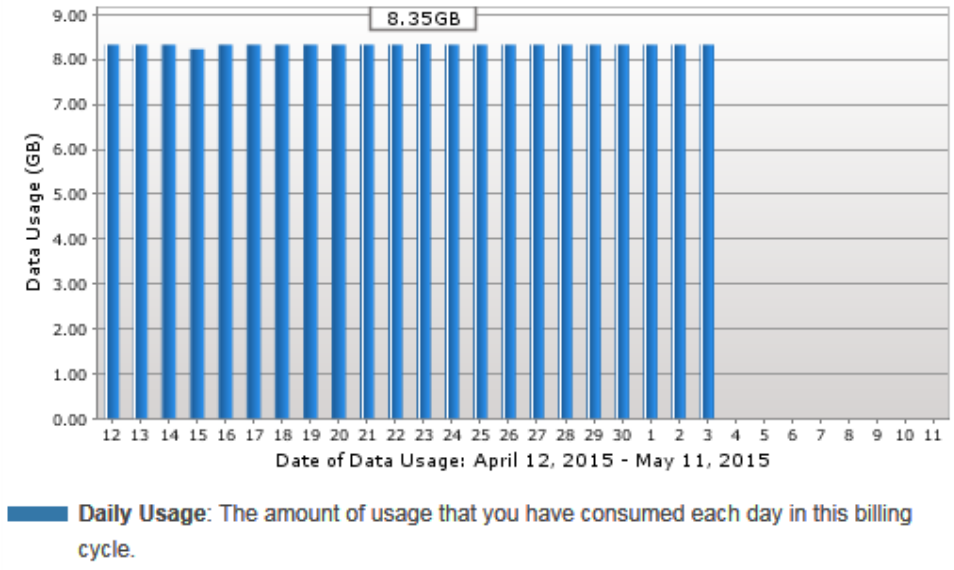

Figure 2 – Sample Graphical Report of Daily Usage

Usage Date	Data Used
04/12/2015	8.33 Gigabytes
04/13/2015	8.33 Gigabytes
04/14/2015	8.33 Gigabytes
04/15/2015	8.23 Gigabytes
04/16/2015	8.33 Gigabytes
04/17/2015	8.33 Gigabytes
04/18/2015	8.33 Gigabytes
04/19/2015	8.33 Gigabytes
04/20/2015	8.33 Gigabytes
04/21/2015	8.33 Gigabytes
04/22/2015	8.33 Gigabytes
04/23/2015	8.35 Gigabytes
04/24/2015	8.33 Gigabytes
04/25/2015	8.33 Gigabytes
04/26/2015	8.33 Gigabytes
04/27/2015	8.33 Gigabytes
04/28/2015	8.33 Gigabytes
04/29/2015	8.33 Gigabytes
04/30/2015	8.33 Gigabytes
05/01/2015	8.33 Gigabytes
05/02/2015	8.33 Gigabytes
05/03/2015	8.33 Gigabytes
Total Usage	**183.18 Gigabytes**

Figure 3 – Sample Table of Daily Usage

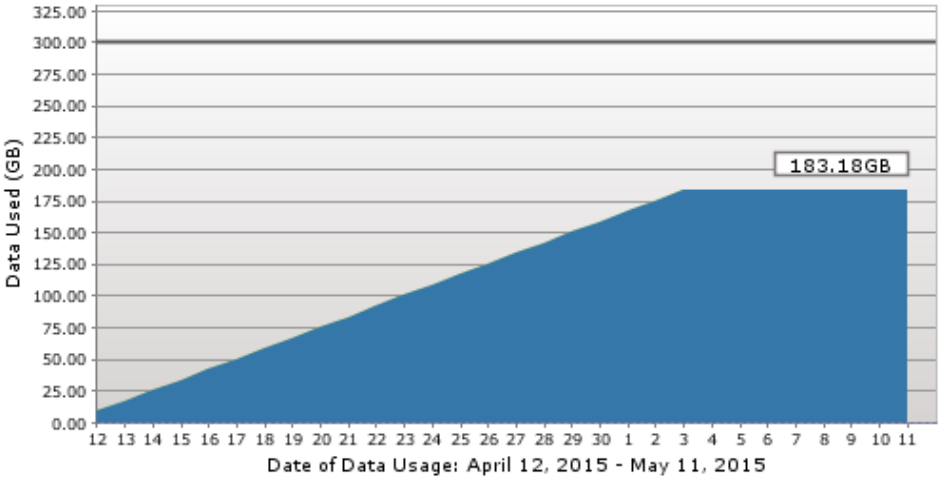

Figure 4 – Sample Cumulative Daily Usage Report

How the Cox Data Usage Meter Works

Cox subscribers connect to the Internet through a cable modem at the subscriber's location, and from there data traffic travels over a local coaxial and Hybrid Fiber-Coaxial (HFC) cable system to a Cable Modem Termination System (CMTS). The traffic then continues through Cox's network and into the Internet.

The CMTS counts downstream and upstream traffic for each subscriber cable modem it serves. Downstream traffic flows from the Internet to the subscriber, and upstream traffic flows from the subscriber to the Internet. The three test sites were supported by a Cisco uBR10000 series CMTS. The CMTS periodically reports the down and upstream counts in an Internet Protocol Detail Record (IPDR) as shown in Figure 5.

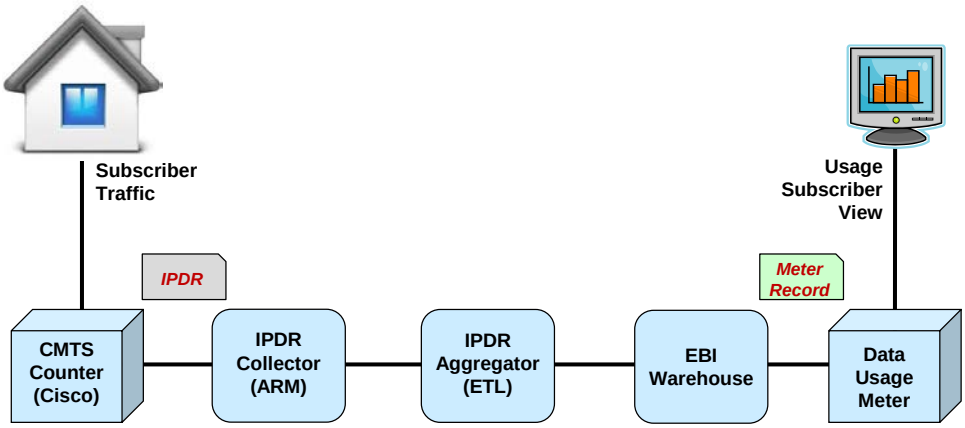

Figure 5 – Usage Meter Process Flow

From the CMTS, IPDRs are sent to the IPDR Collector (Active Resource Manager, or ARM for short), which collects, processes, and stores IPDR data streamed from CMTSs. IPDR processing involves time normalization and usage data mediation, and the ARM also performs checks and cross checks the completeness and correctness of IPDR data to detect anomalies that can result in inaccurate Internet usage data.

In the next stage, the IPDR aggregator gathers data from multiple collectors and converts incremental traffic counts into traffic data in a process referred to as ETL (extract, transform, and load). The IPDR aggregator extracts data from the CMTS's, transforms it to fit operational needs, and loads it into the Cox EBI (enterprise business intelligence) warehouse. Within the EBI warehouse, traffic data is associated with subscriber accounts, and the meter value is calculated for each day to create a meter record. From here the data is fed to the Cox subscriber's data usage meter within the Cox Web portal, which displays the data usage views shown in the previous section.

The requirements for how a cable modem communicates with the CMTS and for how subscriber traffic is transported is defined in the Data over Cable Service Interface Specification (DOCSIS) developed by CableLabs. The IPDR specification is managed by the TeleManagement Forum (TM Forum). A DOCSIS Management Information Base (MIB) defines how traffic is stated in the IPDR.

NetForecast Internet Data Usage Meter Accuracy Validation Methodology

Cox defined specifications for the accuracy and functionality of its data usage meter. NetForecast independently validated the meter relative to Cox's stated goals for: traffic counted, data accumulation period, error bounds, timeliness, exception handling, granularity, mathematical consistency, accessability, availability, and clarity as defined in the meter specification. A comprehensive description of data usage meter specification factors is available in NetForecast's report, *ISP Data Usage Meter Specification Best Practices for MSOs*, reference [1].

As the data usage meter system auditor, NetForecast has no stake in the design or implementation of the Cox data usage meter, nor do we endorse the Cox data usage meter specification goals. The objective of our testing is to assess whether or not the system complies with the meter specification.

NetForecast obtained subscriber accounts and performed usage meter accuracy validation testing for the three test locations. Both downstream and upstream testing were performed under a variety of conditions.

NetForecast instrumented the accounts with a test laptop PC running NetForecast software and a Linux-enabled NETGEAR router with NetForecast's Usage Mapping (UMap) firmware. In addition, we used FTP accounts on various NetForecast servers on the Internet. Figure 6 shows the NetForecast instrumentation. A more detailed description of the NetForecast methodology is available in the NetForecast publication *Validating the Accuracy of ISP Subscriber Traffic Usage Meters*, reference [2].

The NetForecast test involved scripts performing an FTP file transfer from one of our test servers to the laptop. The tests consisted of repeatedly transferring files of varying sizes in complex patterns. These tests were also performed as uploads from the laptop to the server. The script generated a log file documenting the transfer results and capturing detailed timing information for each transfer.

Extreme care was taken to ensure that only test traffic was sent or received through the cable modem. The laptop and router were the only devices connected to the cable modem. The laptop was cleansed of all applications that could generate traffic not needed for the tests—and applications were configured to neither request nor receive any software updates. Finally, remote management of the laptop was scheduled not to occur during testing.

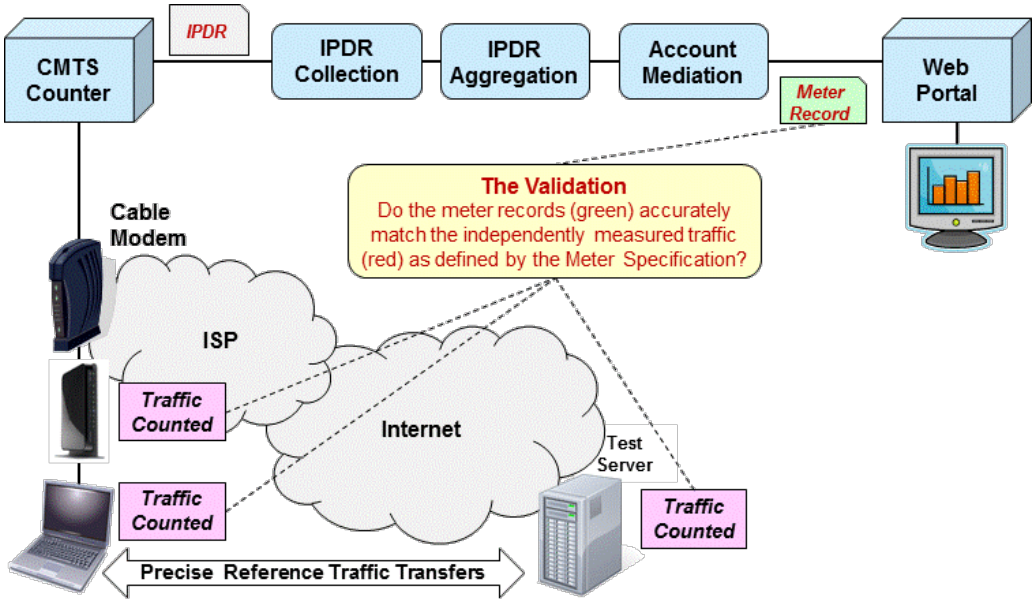

Figure 6 – The NetForecast Methodology

For each test, NetForecast produced three of our own carefully documented records of up and down traffic, illustrated by the pink boxes in the lower left of Figure 6. These record sources included:

- File transfer logs and traffic data on the test PC
- Upstream and downstream traffic processed by the UMap router firmware
- Server records generated using detailed traffic instrumentation software.

Because of their locations in the network, each of the NetForecast measurement points sees different amounts of protocol overhead. NetForecast accounted for upstream and downstream protocol overhead as counted at the CMTS.

NetForecast also gathered data usage meter data from two places in the Cox meter system during testing phase as shown in Figure 6. The data sources were:

- Hourly traffic records from the data warehouse, which provided a preview of measurements presented on the Cox meter portal
- The account portal for the test location, which reported the usage meter data to the subscriber.

Data usage meter system accuracy validation entailed comparing data from the two Cox reporting sources to the three sets of NetForecast-generated measurement data. Comparing the NetForecast and Cox data enabled an end-to-end view of the Cox data usage meter's accuracy.

The hour-by-hour data records collected by NetForecast from the three sources described above were not shared with Cox. Cox provided daily usage meter record files for each of the three sites. Cox did not know which tests NetForecast performed or when they were performed.

NetForecast Independent Cox Data Usage Meter Validation Results

This report section summarizes NetForecast's validation test results for Cox's subscriber Internet data usage meter as seen from the subscriber perspective in the three geographically separate measurement locations. NetForecast validated Cox's specified goals for the data usage meter against what we actually found.

1. What the meter counts

Cox Goal: The Cox data usage meter counts all IP traffic crossing the cable modem-CMTS HFC connection, except for Cox-specific service traffic:

- Video accessed through Cox TV Connect
- Cox Video on Demand and TV accessed through the Contour app and not accessed via a 3rd party app like HBO GO or ESPN3
- Cox Digital Telephone

[Note: NetForecast did not test using Cox-specific service traffic, so we did not validate whether or not the meter counts it.]

The Cox data usage meter also does not count the DOCSIS traffic, which results from communication between the CMTS and the cable modem.

The meter system counts the following traffic classes as subscriber usage:

- Subscriber-generated traffic (also called payload)
- Protocol overhead
- Background traffic (non-subscriber-generated traffic)

[Note: Overhead is contributed by a number of network protocols including: DHCP, DNS, Ethernet, IP, and TCP. These protocols are counted. The overhead varies by application, depending on the protocols the application employs to transmit user data over the Internet. These protocols are necessary for your Internet service to function properly.]

> **NetForecast Validation Result:** NetForecast confirms that all traffic transiting the Ethernet port between the cable modem and home router is counted. This includes overhead contributed by a number of network protocols such as: DHCP, DNS, Ethernet, IP, UDP, and TCP. Additional protocols may be in use based on the devices in the home and applications they operate. DOCSIS traffic was not counted.

2. Data accumulation period

Cox Goal: The meter data accumulation period defines the minimum period over which the meter needs to gather a sufficient number of measurements to be accurate. Cox uses two accuracy timespans: cumulative daily sums, and the monthly total.

> **NetForecast Validation Result:** NetForecast confirms the use of daily cumulative, and one month total data accumulation periods. NetForecast performed its analysis using the same data accumulation time frameworks.

3. Error bounds (Accuracy)

Cox Goal: The error bounds define the acceptable error range goal for the meter. The certification goal for the Cox data usage meter is to be accurate within +/-1% for both the cumulative daily and monthly total values.

> **NetForecast Validation Result:** NetForecast validated that the Cox meter system was accurate to within +/-0.2% at the cumulative daily period and +/-0.1% over the data accumulation period of one month for the sites validated. This includes all the protocol overhead as defined in "What the Meter Counts" above. NetForecast confirms that the monthly data accuracy is well within the Cox goal of +/-1%.

The overall error range across all three tested cities was within +/- 1% over a month. The reporting system maintained that accuracy throughout all the elements of the Cox meter system—up to and including the final view as seen on the customer portal. This statement applies when comparing the meter with the actual traffic sent to and from the cable modem, which includes protocol overhead as described above.

Figure 7 shows the cumulative daily error distribution across the three sites audited. A negative error indicates that the Cox meter value is low relative to the NetForecast reference value (under reporting). A positive error indicates over reporting. All of the tests fell within the plus or minus (+/-) 1% range. The overall mean for all tests was slightly negative, indicating a slight bias toward under reporting.

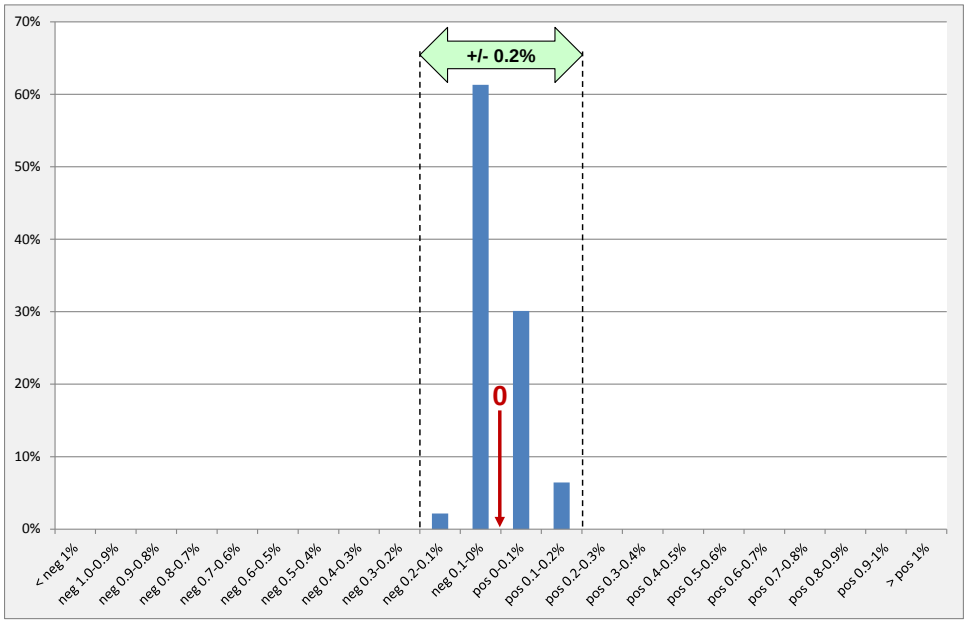

Figure 7 – Cox Cumulative Daily Data Usage Meter Error Distribution

4. Timeliness

Cox Goal: The Cox data usage meter system goal is to gather subscriber usage information hour by hour and roll it up into a 24-hour report by 9AM the following day. A day of traffic is defined as midnight to midnight in the subscriber's local time (adjusted for daylight savings time as appropriate).

> **NetForecast Validation Result:** NetForecast confirms that the customer portal was updated as specified during the test period, except during maintenance periods.

5. Exception handling

Cox Goal: Cox provides alert messages to inform subscribers when delays occur in the display or updating of data.

Message 1: This message appears on the data usage meter page when the data usage meter is undergoing maintenance and the data usage meter is not able to be displayed.

"Your Data Usage Meter is undergoing maintenance. Please try again in 24 hours."

Message 2: This message appears on the data usage meter page when the data usage meter is undergoing maintenance.

"Your Data Usage Meter is undergoing maintenance and your usage may not be current. Please try again in 24 hours."

Message 3: When new subscribers log into the data usage meter page during day one, they will see this message.

"Welcome to Cox High Speed Internet. As a new customer your data usage may take up to 24 hours to display."

> **NetForecast Validation Result:** NetForecast did not experience the error conditions listed above during testing, so cannot confirm that the exception handling messages appear as stated.

6. Granularity (How the data is shown)

Cox Goal: The Cox goal is to present meter data to the subscriber in binary bytes displayed as gigabytes (GB) to two decimal points, using standard mathematical rounding.

> **NetForecast Validation Result:** The Cox goal of presentment to two decimal points, using standard mathematical rounding was met.

7. Mathematical consistency

Cox Goal: The sum of the individual traffic consumption values displayed to subscribers for each day should equal the monthly total displayed to subscribers.

> **NetForecast Validation Result:** NetForecast confirms that the sum of the daily values equals the monthly total displayed to subscribers.

8. Accessibility

Cox Goal: The Cox data usage meter reporting screens will be accessible within one click of user login.

> **NetForecast Validation Result:** NetForecast confirms that the data usage meter is accessible within one click of user login.

9. Availability

Cox Goal: The Cox online data usage meter at the Cox customer portal (http://www.cox.com/) will be available 99% of the time.

> **NetForecast Validation Result:** NetForecast confirms that the customer portal view was available when accessed, but did not continuously monitor to determine if the 99% availability goal was met.

10. Clarity

Cox Goal: The Cox data usage meter will be explained in a FAQ and "learning pages" accessible from the Cox customer portal.

> **NetForecast Validation Result:** NetForecast confirms that the FAQ page and the "learning pages" are clear and informative.

Useful Information if You Want to Do Your Own Testing

If you wish to perform your own Internet usage meter validation testing, you should be aware of factors that may cause your measurements to vary from what the meter reads. One such factor is where you measure. If you measure from the network, you will see protocol overhead that you will not see from a computer. Another factor affecting measurements is packet loss, and another is the presence of "unexpected" traffic.

Avoiding binary versus decimal math confusion

The Cox data usage meter reports 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		
KB	Kilobyte	1,024
MB	Megabyte	1,048,576
10 MB	Megabyte	10,485,760
100 MB	Megabyte	104,857,600
1000 MB	Megabyte	1,048,576,000
GB	Gigabyte	1,073,741,824
TB	Terabyte	1,099,511,627,776

Decimal	
Thousand (Kilo)	1,000
Million	1,000,000
10 Million	10,000,000
100 Million	100,000,000
1000 Million	1,000,000,000
Billion	1,000,000,000
Trillion	1,000,000,000,000

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

- One GB is 2.4% larger than 1000 MB (green vs. pink in the table above). 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.
- 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). This may falsely lead you to believe that the meter is over counting.

For example, if you look at the size of a file on a PC, that value does not include any protocol overhead which may lead you to erroneously conclude that the Cox meter is over reporting.

If you measure traffic in the network, you will see the payload traffic plus overhead from protocols like TCP/IP and Ethernet, which generally add up to 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 your counter should closely track the traffic as seen in the network.

Packet loss adds to the meter

Any reliable end-to-end protocol (like TCP) has a mechanism to retransmit packets lost in transit. Packet loss will add to the traffic seen by the meter in some situations.

"Unexpected" traffic

If you look closely at your Internet data usage, chances are you will see unexpected bytes register on your meter. Most Internet users don't know how much traffic their household produces. The amount of data you consume depends on what you are doing—not the amount of time you spend online. Here is a sampling of traffic sources that may surprise you.

One likely surprise is the number of traffic-generating devices and users in your home. Most subscribers have a wireless (Wi-Fi) router that provides access not just to PCs, but everything from smart phones, tablets, digital video recorders and printers, to game consoles and cameras. Many of these non-PC devices "phone home" to a manufacturer or support service, and for convenience, these automated connections are invisible to the user, so you may be unaware of the traffic generated.

Your internet-connected devices may be infected with malware/spyware, which can generate traffic that has nothing to do with any user in your home. Also, if you fail to secure your home Wi-Fi Internet connection, others may use it without your permission or knowledge, thus increasing your internet traffic.

The most common source of unexpected traffic, however, is PC software. Most popular software has automated update features which download and install updates. This transparent automation is for your convenience and protection, but the traffic it generates may come as a surprise.

Although each program update may be small, but when you multiply a modest download by the number of programs calling for updates and the number of PCs in the house, such traffic can be substantial. Furthermore, in some cases vendor default settings are often

aggressive, checking each hour and downloading every possible option, even if they aren't needed. For example, a software program may automatically load its interface in a dozen languages for a monolingual household.

Another possible “surprise” upstream traffic source is cloud services. These may include such services as online file backup, uploading to photo sharing sites, etc. Again, the software or service settings may be more aggressive than needed.

In addition, many news and information services preload content onto a subscriber's PC, tablet or smart phone over the home Wi-Fi network. The content often arrives overnight for convenient early morning viewing. Of course users don't read all the content every day, but they probably do enjoy the speed with which content appears on the screen. Fresh content may also be sent overnight to a smart phone or MP3 player to be viewed or listened to during the morning commute.

Assume each night's upload is only 1GB, which takes up a modest 1GB on the device's storage—and assume that it never consumes more than 1GB because it overwrites the old content with fresh content each night.

Finally, there may be unexpected traffic to non-PCs. For example, traffic may be going to digital video recorders such as TiVo or streaming boxes like Roku. A user in your home may have rented a movie from Amazon or Netflix, etc.—and you may get more bits than you pay for because many services also preload the start of other movies as well as trailers to make them instantly available should they be called for. As in other situations described above, the traffic is generated for your convenience but without your knowledge.

Most of these traffic sources are low, but some can be unexpectedly high if they aggressively load content. You should 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 at reference [3].

Tracking down rogue traffic

If you cannot account for high traffic volume on the meter and suspect some rogue consumption, we recommend performing a controlled test. Plan for a solid period of time when your home can become "digitally silent" (e.g., overnight or on a weekend when away). At the start of the silent period, turn off all devices that can access the Internet. Make sure, however, to keep the router and cable modem operating.

At the end of the digital silence, turn on one PC and log back into the Cox data usage meter portal, or check from another location while you are away. Look at the daily traffic views for the time period that you observed digital silence. If true digital silence was achieved, the meter should not have incremented by more than 0.1GB in any day. If there is significant traffic shown, then there is certainly some other traffic consumer connected through the router.

Conclusions

The Cox subscriber Internet data usage meter is accurate to within plus or minus (+/-) 0.1 percent (one tenth of one percent) in aggregate over the month, and the cumulative daily values are accurate to within (+/-) 0.2 percent (one fifth of one percent). The reporting system maintains that accuracy throughout all the elements of the meter system—up to and including the final view as seen on the customer portal. These accuracy assessments apply when comparing the meter with the actual traffic sent to and from the cable modem, which includes some protocol overhead as explained above. Based on our test results, subscribers should be able to rely on the meter's accuracy.

NetForecast found the Cox data usage meter online reports easy to understand and use. NetForecast found the online system was always available when accessed during the testing period. We also found the descriptive information about the data usage meter to the site to be adequate and clear.

References

1. Sevcik, *ISP Data Usage Meter Specification Best Practices for MSOs*, NetForecast Report 5110 (Updated), February 2013.

2. Sevcik, *Validating the Accuracy of ISP Subscriber Traffic Usage Meters*, SCTE Cable-Tec Expo, Atlanta, GA, November 15, 2011.

3. 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 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.