



Fourth Accuracy Validation of Comcast's Internet Usage Meter

Peter Sevcik
June 2015

Comcast's data usage meter system measures and reports how much Internet data traffic a subscriber consumes and generates. Comcast has engaged NetForecast to independently and continuously validate the accuracy of its data usage meter system. Our first report was published in December 2009 [1], followed by a second report in May 2010 [2], and a third in May 2014 [3]. This, the fourth report, details the accuracy results from measurements at 55 Comcast service locations for the 16-month period from January 2014 through April 2015. This report is based on more than 3 million measurements over the 16 months.

For Comcast, NetForecast performs both passive measurements of real user traffic in subscriber homes, and active reference tests in which a NetForecast PC and server generate the only traffic on dedicated test lines in subscriber homes. In combination, these two approaches help identify potential meter accuracy issues quickly, and supply data needed to troubleshoot and resolve issues.

For both approaches, NetForecast performs independent traffic measurements, obtains hourly usage meter records from Comcast for each location, and compares NetForecast's measurements with Comcast's records. Results from 50 passive and 5 reference locations are aggregated in this report.

Comcast established a goal for its Internet data usage meter to correctly measure traffic passing through a subscriber's cable modem within plus or minus (+/-) 1% accuracy over a month. **Our analysis validates that Comcast's meter successfully met its goal with an Apdex rating of "Excellent" for daily accuracy and an average 90th percentile month end range of +0.4% overreporting and -4.0% underreporting.** These results mean that the Comcast meter did not over count during the validation period, but that some traffic was not counted.

This report describes how the Comcast usage meter works, the NetForecast meter validation methodology, NetForecast's analytical methodology, NetForecast's meter validation results, and useful information should you wish to perform your own validation of the accuracy of Comcast's usage meter.

The Comcast Usage Meter

Comcast's Internet usage meter provides subscribers with information about how much traffic has traversed their residential Internet connections.

Where subscribers can find their meter report online

Meter reports are available online at Comcast's customer portal, accessible at <http://customer.comcast.com>. After signing in and selecting "My Services," a new page shows various account management tools along with an "Equipment" box. In the Equipment box, the data usage accumulated during the month to the current date is displayed as shown in Figure 1.




Figure 1 – Cumulative Daily Usage View

NetForecast Report
NFR5120

©2015
NetForecast, Inc.

Selecting “View Details” brings up a comprehensive page of information as shown in Figure 2. Subscribers can view usage for the current month as well as a rolling 3-month history.

Selecting “For more information and frequently asked questions” brings up detailed information about the Comcast data usage plan. The page shown in Figure 2 also displays a link to a Data Calculator which subscribers can use to estimate their data usage based on the number of Internet-connected devices and daily Internet activity.



*Actual usage shown is not real-time and could be delayed up to 24 hours.

Usage History

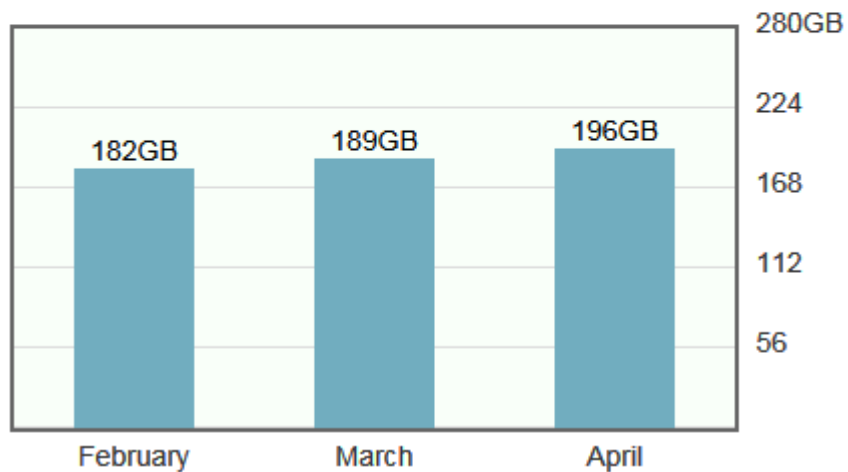


Figure 2 – Month End Usage View

A subscriber can access the customer portal from any browser from any ISP's network, allowing subscribers to check their meter when away from home.

How the meter works

Comcast subscribers connect to the Comcast network through a cable modem and from there 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 Comcast'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. A CMTS periodically reports the downstream and upstream counts in an Internet Protocol Detail Record (IPDR) as shown in Figure 3 (typically every 15 minutes, depending on the CMTS manufacturer).

As Figure 3 shows, Comcast's IPDR system collects and aggregates data from each CMTS and transfers it to the Comcast usage management platform, which associates it with the customer account and stores it in a database. When a subscriber accesses the web

portal, a web service queries the usage management platform database to retrieve the usage data and display it as shown in Figures 1 and 2.




Figure 3 - How the Meter Data Is Processed

The subscriber sees results for devices that are authorized under the subscriber's account. This is typically the cable modem, which is identified by its MAC address. Subscribers with more than one device see a separate meter for each device.

What the meter shows

There are several steps between the time a subscriber's packet moves through the cable modem and when the meter results appear on the Comcast subscriber portal. Each of these steps takes time. We have already described the time lag associated with CMTS traffic reporting—generally every 15 minutes. The IPDR processing system aggregates the traffic and summarizes it by hour. The usage management platform database receives the updates and prepares the results to be shown on an hourly basis. This processing introduces an overall time lag that causes the meter to update about three hours after the subscriber traffic passed over the network. The Comcast goal is to have the meter update no later than 24 hours after the traffic event.

The usage management platform accumulates the upstream and downstream traffic bytes over the month, converts the total to gigabytes (GB), and truncates to whole gigabytes. Truncation means that a unit value is not shown until a full decimal value is accumulated; e.g., 9.9 is truncated to 9, and 10.1 is truncated to 10, etc. The result displayed in the portal is, therefore, the cumulative whole GB sum of all traffic from the beginning of the month (down + up). At the start of each month the meter resets to zero but displays "<1GB."

The meter operates on Coordinated Universal Time (UTC, also known as GMT). The "new month" meter reset occurs at 00:00 UTC on the end of the last day of each month, so the change appears during the evening of the last day of the month across the US. For example, in the Eastern Time Zone the month change will occur at 8:00 PM during local standard time (EST) and 7:00 PM during daylight savings time (EDT).

NetForecast Internet Usage Meter Accuracy Validation Methodology

All measurements were performed using the NetForecast UMap service delivery platform shown in Figure 4. The UMap system is enabled by measurement and reporting software embedded into customized, fully-featured, wireless home routers supplied and supported by NetForecast. Once installed, the routers register with the NetForecast UMap management and reporting system.




Figure 4 – NetForecast UMap Service Delivery Platform

UMap is a parallel data usage meter system deployed in subscribers' homes. The UMap platform has a high degree of security and resiliency, and strong safeguards to protect the privacy of participants.

The UMap system operating on the Comcast network performs two types of measurements: active and passive as shown in Figure 5. This report is based on results from 5 active and 50 passive measurement sites.

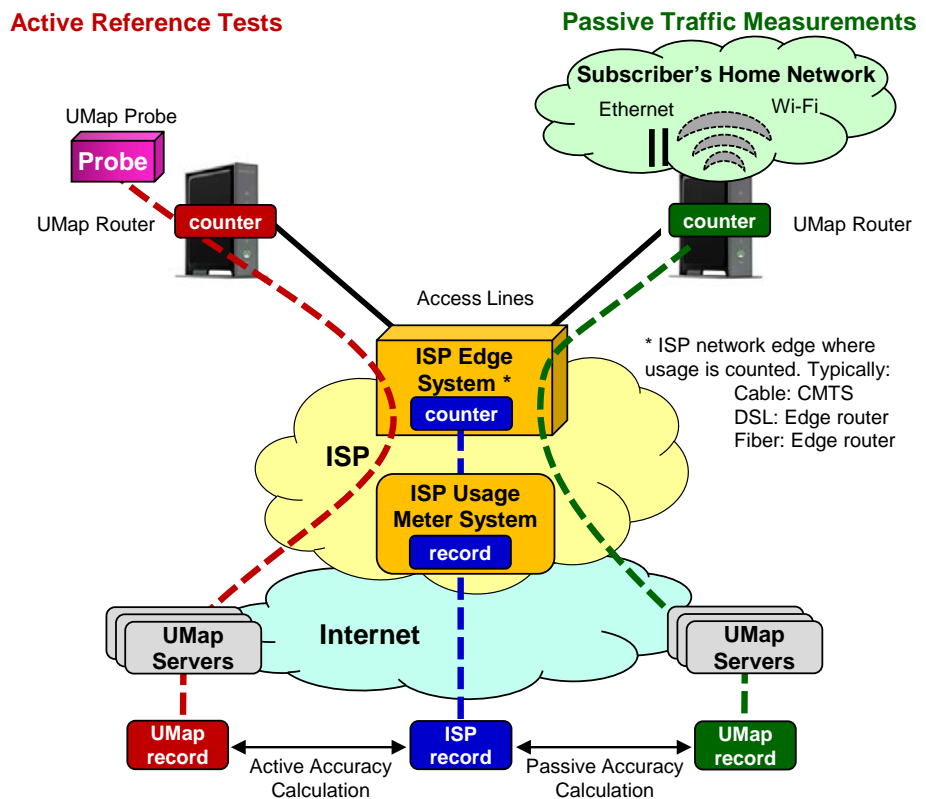


Figure 5 – The NetForecast Methodology

Active reference testing

This description applies to the red dashed line of data flow in Figure 5.

For the active reference test locations, NetForecast installs a test PC running NetForecast software on each Internet connection. We use FTP accounts on NetForecast servers on the Internet. NetForecast obtains subscriber accounts and performs usage meter accuracy validation testing for all active reference test locations. Both downstream and upstream testing are performed under a variety of conditions.

The tests consist of repeatedly transferring files of varying sizes in complex patterns. The system generates a log file documenting the transfer results and capturing detailed timing information for each transfer.

Passive traffic measurements

This description applies to the green dashed line of data flow in Figure 5.

For the passive locations, the UMap system continuously measures real-user traffic traversing the home Internet connections of many actual subscribers. Passive measurements take place under real-world conditions—i.e., it relies exclusively on the subscriber's home traffic. This enables measurements to be performed at many locations, thus providing comprehensive geographical coverage. Unlike active testing, which is not continuous, passive traffic measurement provides data for every hour during the month to give additional insight into meter accuracy.

Calculating Meter Error

UMap data from either the active or passive approach is adjusted to assure that UMap measurements count the same overhead as the CMTS counts. NetForecast then aligns the hourly usage records from UMap and Usage Management Platform so that the same hours are compared. Daily sums are generated for each site. NetForecast applies the following formula to the UMap and the Comcast Usage Management Platform daily traffic measurement records.

$$\mathbf{Error} = \frac{(\mathbf{Comcast\ Record} - \mathbf{UMap\ Record})}{\mathbf{UMap\ Record}}$$

If the error result is positive, the meter is overreporting. If the error is zero, the meter is as accurate as it can be, and if the error result is negative, the meter is underreporting. Results are shown as a percentage. Each site-day error result is an error sample.

Meter Accuracy Validation Findings

Comcast worked with NetForecast to create a Comcast-specific Internet Data Usage Meter Accuracy Specification that defines goals for factors against which NetForecast performed its independent meter accuracy validation. The following table summarizes the goals for the specification factors, and Comcast's performance relative to those goals. The first 4 specification factors relate to measuring meter accuracy, and the remaining 6 relate to assessing the meter user experience via the subscriber portal. A comprehensive description of data usage meter specification factors is available in NetForecast's report, *General ISP Data Usage Meter Specification and Best Practices*, reference [4]

Meter Factor	Comcast Goal	NetForecast Assessment
Meter Accuracy Factors		
What Is Counted	Count all subscriber-generated IP traffic across the subscriber's Internet access line, including IP protocol management traffic and Ethernet framing.	NetForecast validates that the meter counts as specified.
Meter Record Update Rate	Aggregate, mediate, and store data as a formal meter record on an hourly basis (UTC time).	NetForecast validates that the meter records are processed and stored on an hourly basis.
Accuracy Time Period *	Measure accuracy on a cumulative daily and month end basis.	NetForecast measured accuracy on a cumulative daily and month end basis. *
Error Bounds *	Meet an accuracy goal of +/-1% on a cumulative daily and month end basis.	NetForecast validates that the meter met the daily accuracy goal with an overall Apdex score of Excellent over the 16 months. NetForecast also validates that the month end error range during this period had an average overreporting error of +0.4% and underreporting error of -4.0%*
Meter Portal Experience Factors		
Timeliness	Update the subscriber portal 3 to 24 hours after the end of a usage event	NetForecast confirms that the subscriber portal updated within 3 to 24 hours after the end of a usage hour.
Granularity	Present usage to the subscriber in binary bytes displayed as GBs truncated to the whole GB.	NetForecast confirms that the meter data is presented in binary bytes displayed as GBs truncated to the whole GB was met.
Mathematical Consistency	Sum of individual values shown for each display period within a month match the total value shown for the month.	This goal does not apply as presently there are no subtotal values shown on the subscriber portal.
Accessibility	Meter data on the portal is accessible within 3 clicks of subscriber login.	NetForecast confirms that the goal of accessibility within 3 clicks of subscriber login was satisfied.
Availability	The online data usage meter is available 99% of the time.	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.
Clarity	The meter is adequately explained in a FAQ accessible from the subscriber portal.	NetForecast confirms that the FAQ page is clear and informative consistent with Comcast's goal.

* See next section for more details.

Comcast Meter Accuracy Assessment

This section provides more detailed findings regarding Comcast's accuracy performance history relative to its error bounds goals for both the month end and cumulative daily accuracy time periods. NetForecast performed its detailed accuracy analysis by: determining the 90th percentile of sites based on month end error data; then determining the daily accuracy of those sites for each month; and finally aggregating and summarizing the results.

Cumulative Daily Error

A subscriber should be able to use the meter portal to track and manage usage. The only tool is the cumulative daily usage view shown in Figure 1. The cumulative daily error should stay within +/-1%. Determining the effective accuracy for this value across 55 sites over 30 days is complex. Figure 6 uses the Apdex methodology to show how well Comcast met the goal over the last 16 months (see Appendix A for more information). An Apdex score Excellent is desirable and a Good score is acceptable. Comcast achieved and Excellent or Good score in all months. The average Apdex score over the 16 months is 0.94 or Excellent, thus Comcast met the overall objective of informing subscribers with accurate usage meter data.

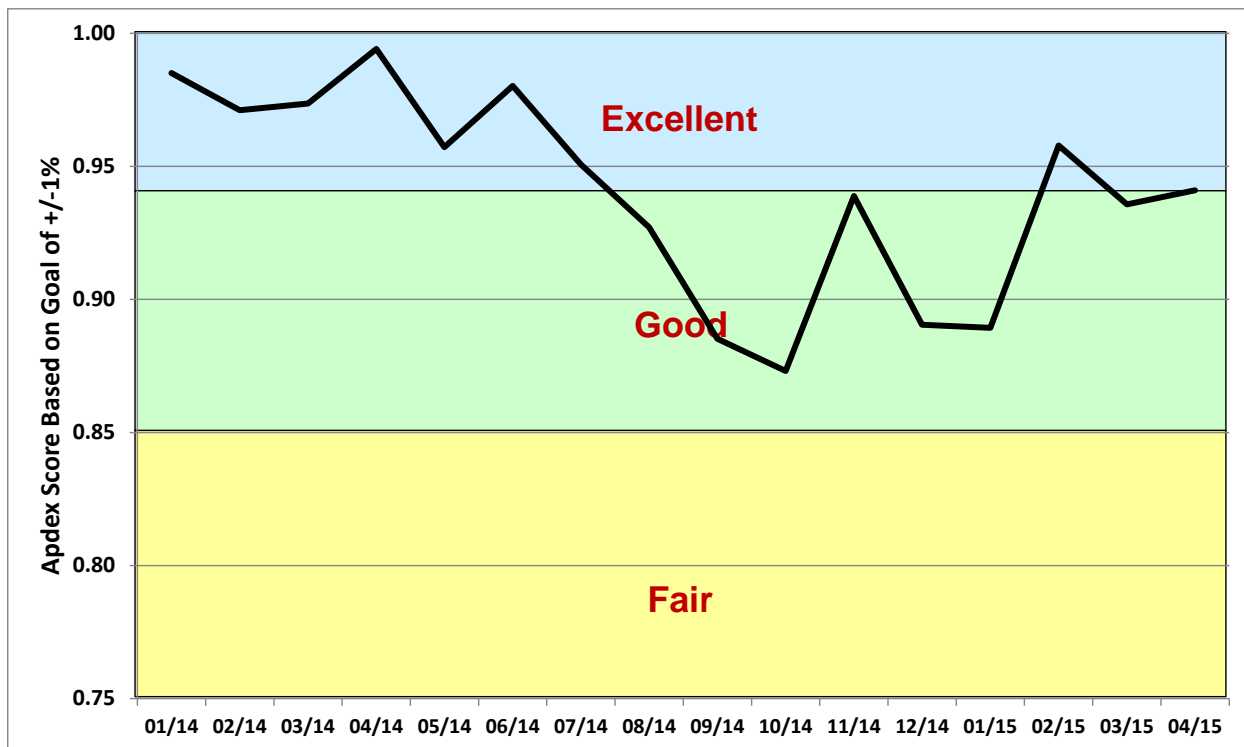


Figure 6 – Daily Meter Accuracy

As Figure 6 shows, however, not all months had Excellent Apdex scores. From August 2014 through January 2015 Comcast's Apdex score was in the Good zone. During these months the underreporting value of the month end error range also dipped below -1.0%.

When the Apdex score began to drop, Comcast investigated, identified, and corrected the cause, after which Apdex scores rose. The root cause was that early in 2014 Comcast upgraded to a new, more scalable IPDR collection platform. Hardware limitations in the new platform caused sporadic underreporting that began in August 2014 and peaked in

October, 2014. From January through early February, 2015 Comcast upgraded the platform hardware, which successfully eliminated the capacity problems.

In another view of Comcast’s meter accuracy, Figure 7 shows a histogram of all cumulative daily errors across all sites within the month end error 90th percentile over 16 months. The figure shows how cumulative daily error rates map to Apdex results. Note that the distribution is generally normal around zero with the some skew towards underreporting evident in monthly chart in Figure 6.

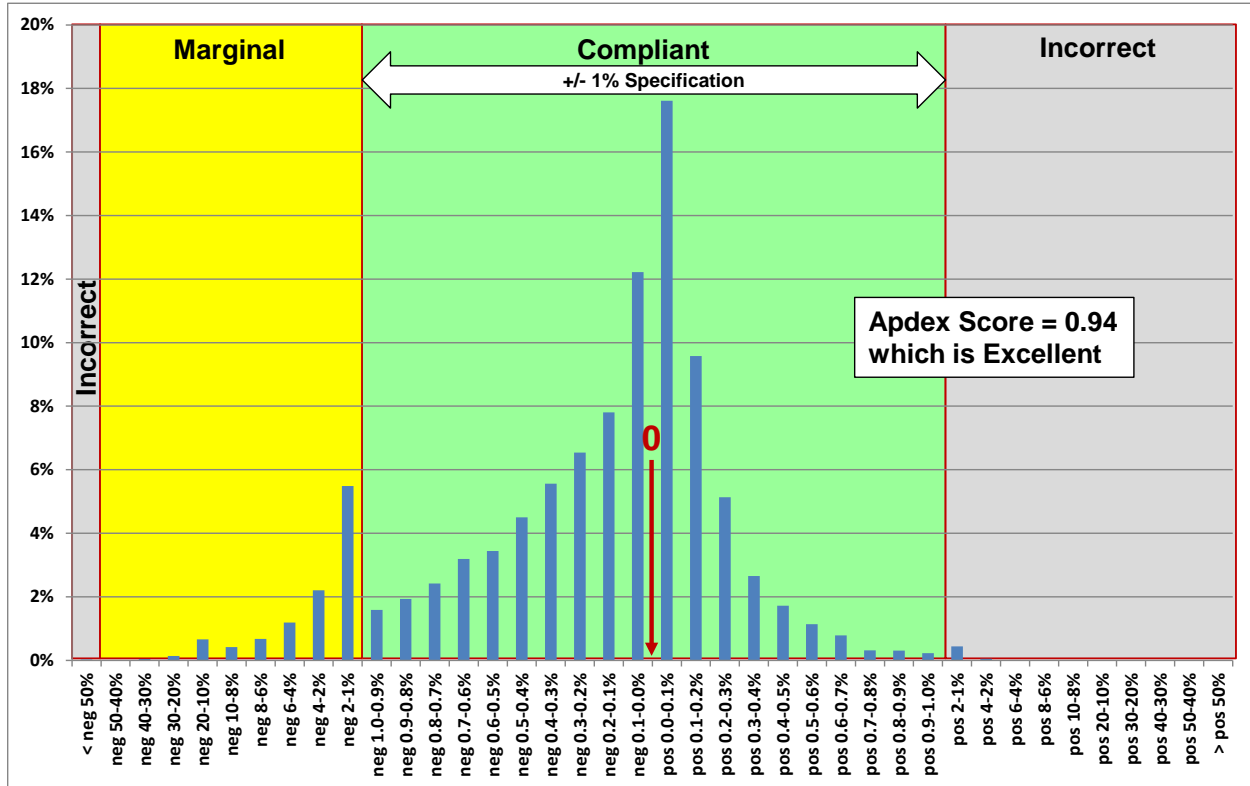


Figure 7 – All Cumulative Daily Results Mapped to Apdex Categories

Month End Error

NetForecast used the month end error range analysis described in Appendix A to determine the 90th percentile boundaries of the month end values for each month. In general 55 sites supplied results into a monthly report. The monthly total of active sites varies due to volunteer churn. Given a hypothetical month during which 48 actives sites supplied sufficient days of measurement data, the process described in Appendix A calls for removing the top 5% (most overreporting) and bottom 5% (most underreporting) sites. The formula $(0.05 \times 48 = 2.4)$ yields 2.4 which is rounded to 2. Therefore in this example the 2 most overreporting sites and 2 most underreporting sites are removed. The next most over and underreporting sites define the month end error range.

The month end error range never exceeded the critical +1.0% goal during the 16-month measurement period. Some months had values below the -1.0% goal. These modest underreporting months track the lower Apdex values of the daily errors shown in Figure 6. The average of all underreporting error range values over the 16 months was -4.0%. Such modest underreporting is acceptable since it favors the subscriber, but is not so low that it undermines trust in the system.

Useful Information if You Want to Do Your Own Testing

If you wish to perform your own Internet Usage meter validation testing, it is important to be aware of factors that may cause your measurements to vary from what the meter states. 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 Comcast 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			Decimal	
KB	Kilobyte	1,024	Thousand (Kilo)	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
GB	Gigabyte	1,073,741,824	Billion	1,000,000,000
TB	Terabyte	1,099,511,627,776	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).

For example, if one looks at the size of a file on a PC, that value does not include any protocol overhead which may lead the user to erroneously conclude that the Comcast 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 it should closely match the network view.

Packet loss affects 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 and subtract in other situations.

“Unexpected” traffic

If you look closely at your Internet usage, chances are you will see unexpected bytes register on your meter. Most Internet users don't know how much traffic their household produces. Here is a sampling of traffic sources that may surprise you.

Peer-to-peer (P2P) file sharing applications like BitTorrent, often used to exchange music, movie, and other types of large files, are a common source of unexpected traffic. Typically P2P agents operate automatically in the background on your network exchanging large amounts of data over your Internet connection. Users are often unaware of the total traffic consumed by P2P over a month.

NetForecast's research has recently uncovered unexpected high consumption caused by cloud storage and file sharing services such as Google Drive, OneDrive, Box, Dropbox, and iCloud. A set of conditions can cause some cloud services to 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 very high data consumption, we suggest that you investigate your configuration with the help of your cloud service provider.

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 to 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.

In rare cases, a PC could be hijacked and generate traffic that has nothing to do with any user in your home. Also, a neighbor or “wardriver” may use your connection without your permission or knowledge.

The most likely 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 download may be small, 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 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 traffic source is 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 or tablet. 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 tablet 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. As modest as this may seem from a device storage point of view, that 1GB did consume bandwidth each night, adding up to 30GB over a month on the meter (plus protocol overhead).

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 the home may have rented a movie using a subscription to a steaming service—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 [5].

Tracking down rogue traffic

If you cannot account for a 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 the home can become “digitally silent” (e.g., overnight or on a weekend when traveling). At the start of the silent period, record the data usage displayed on the Comcast meter, then 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 computer and log back into the Comcast usage meter portal, or check from another network. If true digital silence was achieved, the meter should not have incremented at all during any day. Because Comcast rounds down, values may increment to the next GB if the previous value was close to the rounding threshold before the silent period began. Subsequently, however, it should remain at that value for the rest of the digital silence period. If there is significant traffic shown, then there is certainly some other traffic consumer connected through the router.

NetForecast provides deep insights into difficult network performance problems.

Additional information is available at:
www.netforecast.com

NetForecast and the curve-on-grid logo are registered trademarks of NetForecast, Inc.

Conclusions

Comcast established an accuracy goal for its Internet data usage meter to correctly measure traffic passing through a subscriber's cable modem within +/-1% over a month. Based on hourly measurements at 55 measurement sites from January 2014 through April 2015, NetForecast validates that the Comcast Internet data usage meter cumulative daily accuracy achieves an Apdex score of Excellent. The average month end error range was +0.4% to -4.0%.

Based on our measurement results, subscribers should be able to rely on the meter's accuracy, with a modest trend toward under counting. The Comcast meter also met specified goals for timeliness, granularity, accessibility, availability and clarity during the 16-month validation period.

References

1. Sevcik, [Comcast Usage Meter Accuracy](#) NetForecast Report 5101, December 2009.
2. Sevcik, [Comcast Usage Meter Accuracy \(Updated\)](#) NetForecast Report 5101, May 2010.
3. Sevcik, [Third Accuracy Assessment of Comcast's Data Usage Meter](#) NetForecast Report 5116, May 2014.
4. Sevcik, Wetzel, and Lacy, [General ISP Data Usage Meter Specification and Best Practices](#), NetForecast Report 5119, May 2015.
5. 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 three 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.

APPENDIX A – How to Report Meter Accuracy

The Comcast meter portal view, and example of which is shown in Figure 2 presents two views of usage to subscribers:

Cumulative Daily Usage View: The sum of usage by day from month start to month end. The usage bar in the top of Figure 2 is analogous to an automobile gas gauge. It shows how much “fuel” (Internet capacity) has been used. Subscribers can use this gauge to manage data usage over the month. In a typical month, 30 samples per site can be used for error evaluation.

Month End Usage View: The total usage for the complete month as shown at the bottom of Figure 2. This value is the same as the last value in the cumulative daily usage view. Subscribers see the total volume of consumption, which provides a historic view of usage. The month end view provides only one sample per site for error evaluation.

NetForecast compares its independently measured traffic counts of synthetic and/or real user traffic with the counts generated by Comcast’s meter system for the same traffic, to determine whether the comparisons fall within Comcast’s target accuracy range.


Given differences in the nature and sample size of the two Comcast views, NetForecast performs different error analyses for each view as described below. NetForecast’s analysis objective is to evaluate the accuracy of the two views from the perspective of *what matters to the subscriber*.

Cumulative Daily Error Analysis

Background: Analyzing and reporting on the error of a system is complex. One could take a simple approach of averaging the error of all samples. If the system has a significant bias, the simple average would show that bias (e.g., most samples are 10% low). However, if a system is fundamentally accurate, the mean (average) or median will be essentially zero. The median of the Comcast cumulative daily error samples is within +/-0.1%. But that result tells us that half of the samples are higher and half are lower. Many samples may be far from the median; therefore, usage meter accuracy assessment should not use averages.

Although under most circumstances subscribers receive accurate meter information, it is not uncommon for a meter system to occasionally provide inaccurate meter information affecting only a few subscribers. These infrequent events are typically called the “long tail of a statistical distribution.” This may seem unimportant, but if the long tail encompasses a large sample size, the meter cannot be considered accurate.

Assessing Accuracy from the User’s Perspective with Apdex: To provide clear insight into accuracy, NetForecast applies the Accuracy Performance Index (Apdex)—an open standard that defines a method for reporting the meaning of many measurement samples from the user’s perspective. Apdex provides a uniform way to analyze and report on the degree to which measured accuracy meets a specific goal.



Apdex is a way to convert many values into a simple meaningful number that properly reflects the user's perspective of performance achieved relative to a performance target.

It is a methodology promoted by vendors and enterprises who are members of the Apdex Alliance. The Alliance uses the rating scale above to communicate scores indicating excellent to unacceptable performance.

Error samples are placed into the following categories each month:

- Compliant:* Number of samples within the +/-1% meter specification. These samples clearly meet the goal.
- Marginal:* Number of samples between -1% and -50% (underreporting). Underreporting is outside the specification; however, some modest underreporting can be tolerated since the subscriber is not harmed by some traffic not being counted.
- Incorrect:* Number of samples below -50% (excessive underreporting). Excessive underreporting (errors more than -50%) are not tolerated by the user. Excessive underreporting makes the meter useless (think of it in terms of a gas gauge underreporting by 50%) and undermines the subscriber's trust in the meter.
- Incorrect:* Number of samples above +1% (any overreporting). Any overreporting above the 1% specification is a serious problem that users will not tolerate.

The Apdex methodology converts many measurements into one number on a uniform scale of 0 to 1 (0 = completely inaccurate; 1 = perfect accuracy). The resulting Apdex score is a numerical measure of accuracy performance. The Apdex formula is the number of compliant results, plus half of the number of marginal results, plus none of the incorrect results, divided by the total number of samples:

$$\text{Apdex [T]} = \frac{\text{Compliant} + \frac{\text{Marginal}}{2}}{\text{Total samples}} = \text{Score (0 to 1) [T]}$$

Incorrect samples have zero value

Target T or specification is always shown with the score

The value "T" is the target accuracy range of (+/-) 1% (the Comcast goal) and is always shown with an Apdex score.

Month End Error Analysis

Background: There are only 55 month end values per month (one per site). As with cumulative daily results, averaging is not informative. It is important to examine the two edges of accuracy—how far was negative error from zero and how far was positive error from zero?

Simply stating the absolute greatest deviation from zero can lead to extreme results that are misleading outliers. For example, if all but one of the samples are within +/-1% and the one outlier is at +5%, then the maximum of +5% does not accurately represent the general edge of positive error. Trying to apply standard deviation analysis is difficult because there are too few samples and their distribution is not normal.

Error Range Percentiles: The month end error samples are ranked from most positive to most negative as shown in Figure 8. The top 5% and the bottom 5% of the samples by rank are identified. These two samples define the edges of the majority of samples. The error values of these samples are the worst underreporting and overreporting of 90% of the dataset. This is the error range within which 90% of the month end site-errors fall.

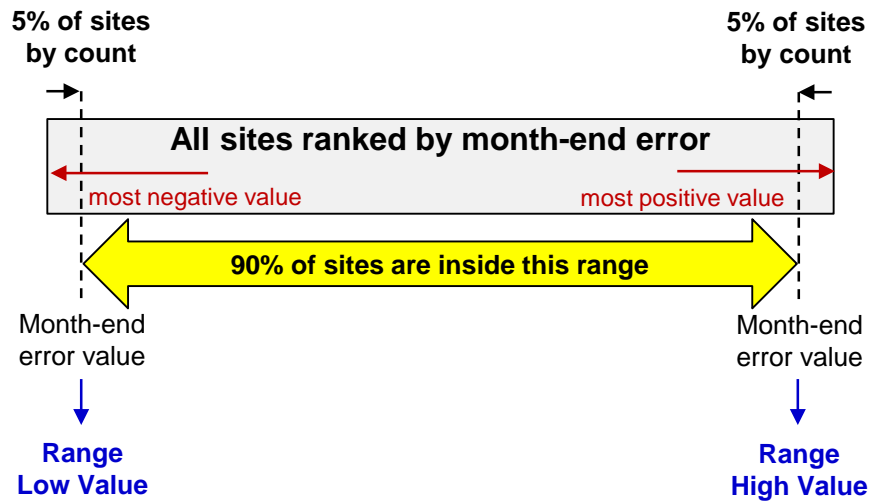


Figure 8 – Calculating the Month End Error Range