Comcast Internet Usage Meter 2019 Accuracy Audit Report

Andrew Lacy, Peter Sevcik, Rebecca Wetzel
June 2020

Executive Summary

NetForecast has independently and continuously audited the Comcast Internet data usage meter since 2009. This 9th report continues that process. Comcast again engaged NetForecast to validate the accuracy of its data usage meter system in 2019. The Comcast meter accuracy specification states that the meter should correctly measure traffic within plus-or-minus (+/-) 1% on a month-end and daily basis. During 2019, Comcast's meter rarely overcounted but sometimes undercounted, causing some user traffic not to register on the meter.

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. 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.

NetForecast gathers usage data throughout the year. The accuracy assessment in this report is based on 3,434,166 audit measurements. NetForecast analyzes each site's data to determine how accurately the meter counted that site's data for the month. The results are reported using three metrics: site monthly error, month-end error (MEE) frequency, and yearly Apdex score.

NetForecast analyzed the number and magnitude of out-of-spec month-end errors normalized across all measured sites and all months in the year using a standard 95% confidence interval calculation. NetForecast validates with 95% statistical confidence that in 2019, across all measured sites, out-of-spec monthly errors had an overall impact of undercounting by -4.74% as Figure 1 shows (-4.92 out-of-spec low plus +0.18 out-of-spec high).

The extent of resulting high (overcounting) or low (undercounting) error conditions applied to all site-months is shown in Figure 2.

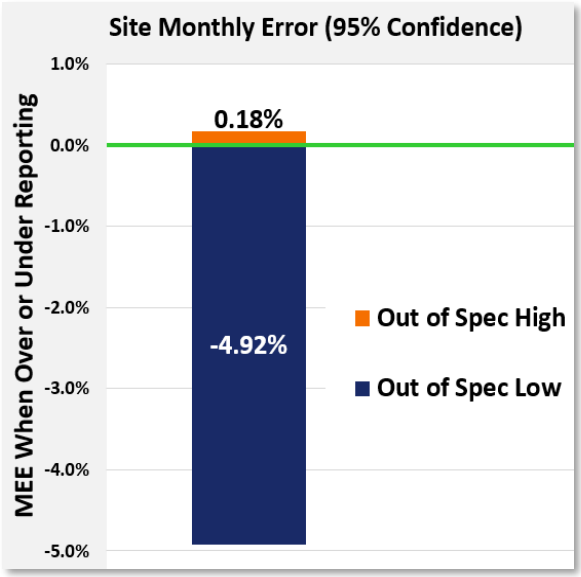

Figure 1 – Site Monthly Error in 2019

Month-end meter errors resulted in an overall Apdex score of 0.97, a rating of "Excellent" for the year (see Appendix A for information on Apdex).

NetForecast's audit methodology has been assessed by NERA Economic Consulting [1], which concluded that the measurement technology is unbiased, and the analytical process is reliable, with appropriate confidence intervals.

NetForecast Report NFR5136

©2020 NetForecast, Inc.

NetForecast Meter Accuracy Assessment

How often and by how much was the meter out-of-spec at month end?

NetForecast determined the frequency and magnitude of month-end errors normalized across all sites and all months using a standard 95% confidence interval calculation. The magnitude of over/undercounting is reflected in the monthly error calculations, which quantify how much the meter would be affected (with 95% confidence) when a site over or undercounted. Only out-of-spec site months (those with values greater than +1.0% and less than -1.0%) are used in this metric.

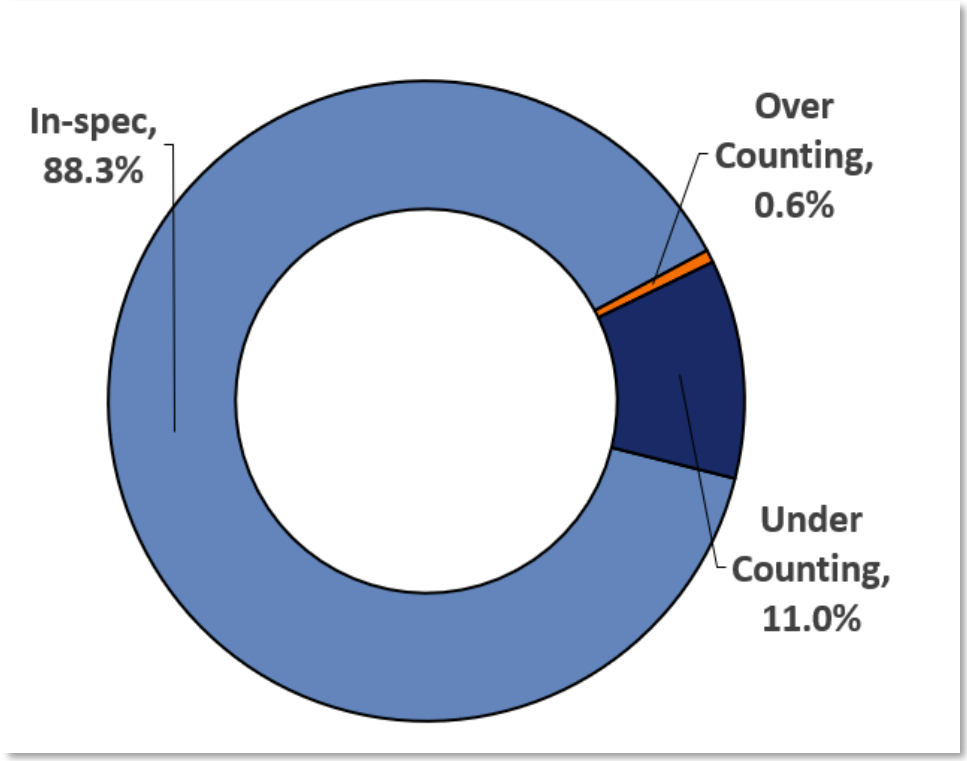

Figure 2 – Comcast Meter Month-end Error Frequency Distribution in 2019

During 2019, the Comcast meter met the accuracy goal at month end 88.3% of the time as Figure 2 shows. Given the margin of error, the sample size, and the standard deviation, the Comcast meter could be expected to overcount 0.6% of the time and undercount 11% of the time in 2019.

As Figure 1 shows, when a site undercounted, the expected underreporting amount was -4.92%, and when a site overcounted, the expected overreporting amount was +0.18%.

Undercounting is most often caused by missed or lost traffic counts. Processing errors or resource limitations in the meter system sometimes result in counts not accumulating in the meter. Overcounting is rare compared to undercounting and is generally caused by retransmissions due to packet loss.

How well did the meter perform on an annual basis?

NetForecast applied the Apdex methodology (see Appendix A) to month-end error results to determine the meter's overall monthly ***accuracy score*** and associated ***accuracy rating,*** also known as the ***yearly Apdex score***. Think of this process as calculating the numeric score on a test that results in a letter grade. For example, an 88% *score* on a test results in a grade (*rating*) of "B." A high score indicates that more site months were within the specification goal of +/-1%. Site month errors above +1% negatively impact the scores much more than errors below -1%.

NetForecast measured and documented the month-end error of Comcast's data usage meter system for every measurement site in every month in 2019. Figure 3 shows the Apdex score history for all years the Comcast meter has been continuously measured with the current NetForecast methodology. Figure 3 shows the Apdex score history.

With an Apdex score of for .97 for 2019, Comcast's data usage meter accuracy remained within the 'Excellent" category for the sixth year in a row.

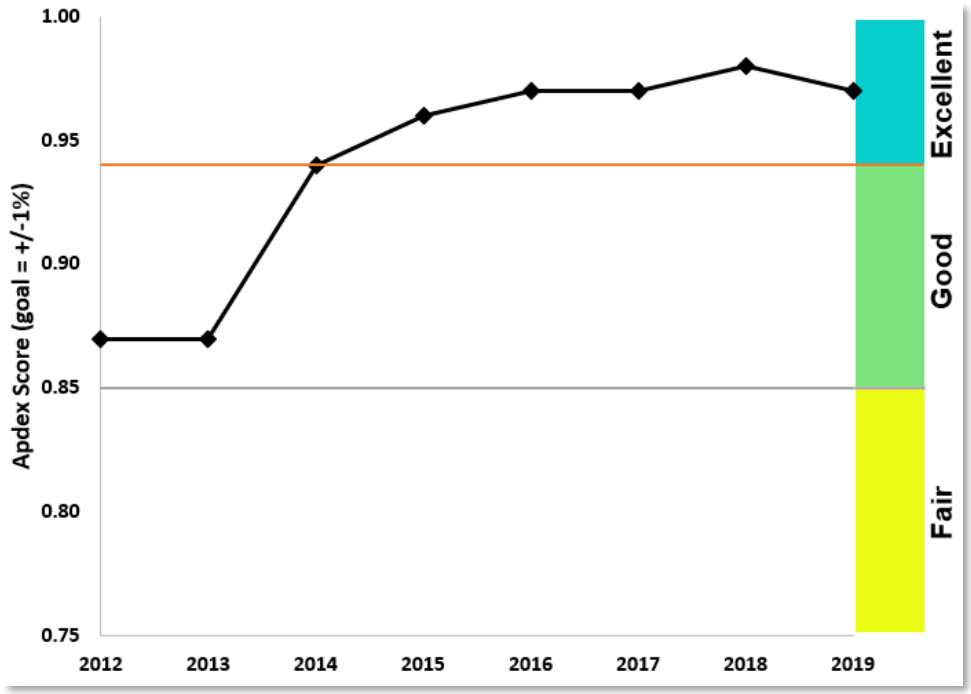

Figure 3 – 2012 through 2019 Yearly Apdex Scores

Meter Accuracy Factors

Comcast worked with NetForecast to create a Comcast-specific Internet Data Usage Meter Accuracy Specification that defines goals which NetForecast assessed. The following table summarizes the goals for the accuracy specification factors, and Comcast's performance relative to those goals. A comprehensive description of data usage meter specification factors is available in NetForecast's report, *General ISP Data Usage Meter Specification and Best Practices* [2].

Meter Factor	Comcast Goal	NetForecast Assessment	Goal Met
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.	Yes
Meter Record Update Rate	Aggregate, mediate, and store data as a formal meter record on an hourly basis.	NetForecast validates that the meter records are processed and stored on an hourly basis.	Yes
Accuracy Time Period	Measure accuracy on a cumulative daily and month-end basis.	NetForecast measured accuracy on a cumulative daily and month-end basis.	Yes
Error Bounds	Meet an accuracy goal of +/- 1% on the overall site monthly error.	NetForecast validates that the meter met the +1% but not the -1% accuracy goal.	+1% Yes -1% No

The Comcast Usage Meter

Comcast's Internet data usage meter provides subscribers with information about how much traffic has passed over 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.xfinity.com. After signing in, select the "Services" tab, then click on "Manage Internet" and scroll to the bottom of the page to see a Monthly Data Usage chart shown in Figure 4. Subscribers can view usage for the current month as well as a rolling 6-month history. Positioning the cursor over any month shows the exact usage value in gigabytes (GB).

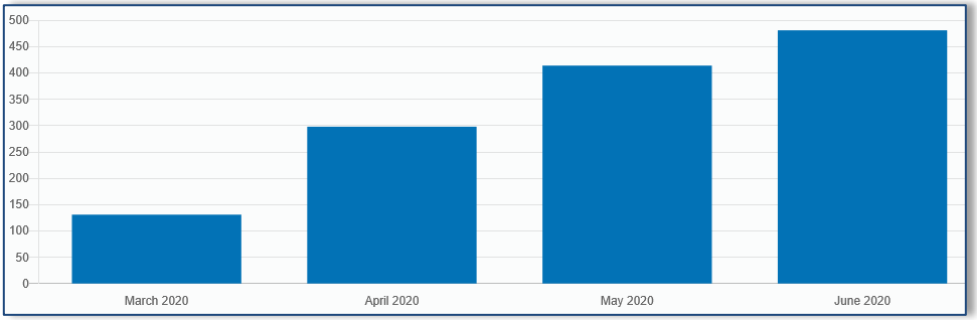

Figure 4 – Cumulative Monthly Usage View

Selecting “Show Table Data” brings up information on the current month and the previous 6 months as shown in Figure 5.

Month	Total Monthly Usage
December 12/01/2019 - 12/31/2019	476 GB
January 01/01/2020 - 01/31/2020	286 GB
February 02/01/2020 - 02/29/2020	192 GB
March 03/01/2020 - 03/31/2020	131 GB
April 04/01/2020 - 04/30/2020	298 GB
May 05/01/2020 - 05/31/2020	414 GB
June 06/01/2020 - 06/30/2020	481 GB

Figure 5 – Tabular 6-Month Usage History

How the meter works

Comcast subscribers connect to the Comcast network through a cable modem, and from there traffic travels over a Hybrid Fiber Coaxial (HFC) cable system to a Cable Modem Termination System (CMTS). The traffic then continues through Comcast's network and to 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 reports the downstream and upstream counts in an Internet Protocol Detail Record (IPDR) as Figure 6 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 4 and 5.

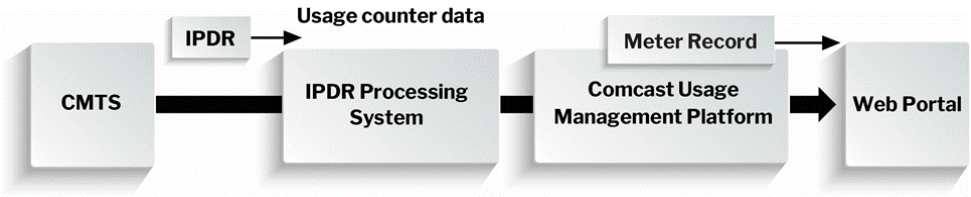

Figure 6 - How the Meter Data Is Processed

What the meter shows

Several steps occur 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. The time lag associated with CMTS traffic reporting is generally 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 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 7:00 PM during local standard time (EST) and 8:00 PM during daylight savings time (EDT).

NetForecast Internet Usage Meter Accuracy Validation Methodology

All measurements were performed using the NetForecast UMap℠ service delivery platform. The UMap system is enabled by measurement and reporting software embedded into customized, fully featured, wireless home routers supplied and supported by NetForecast. The UMap system operating on the Comcast network performs two types of measurements, active and passive, as shown in Figure 7.

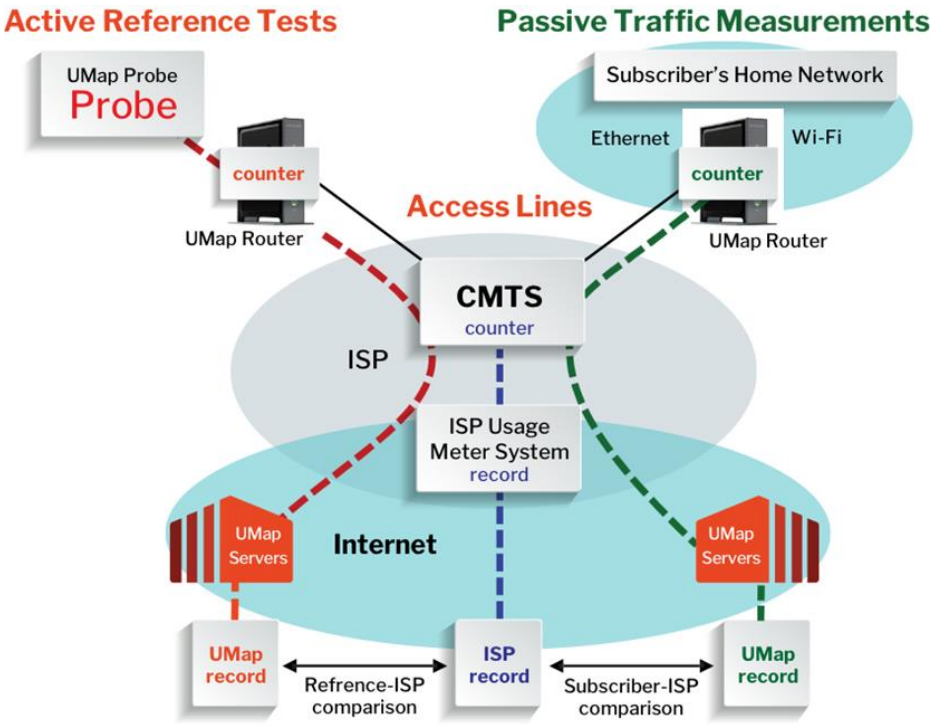

Figure 7 – The NetForecast Methodology

Active Reference Testing

This description applies to the red dashed line of data flow in Figure 7. For the active reference test locations, NetForecast installs a probe running NetForecast software that generates traffic following defined usage profiles under a variety of conditions.

Passive Traffic Measurements

This description applies to the green dashed line of data flow in Figure 7. For the passive locations, the UMap system continuously measures real-user traffic traversing the home Internet connections of many actual subscribers. Passive measurements are made under real-world conditions—i.e., passive measurement relies exclusively on the subscriber's home traffic.

Calculating Meter Error

UMap data is adjusted to ensure that UMap measurements count the same protocol overhead as the CMTS counts. NetForecast then aligns the hourly usage records from the UMap system with the hourly records from the Comcast meter system so the same hours are compared. Daily sums are generated for each site. NetForecast applies the following formula to the UMap and Comcast usage meter daily traffic measurement records:

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

If the error result is positive, the meter is over counting. If the error is zero, the meter is as accurate as it can be, and if the error result is negative, the meter is under counting. Results are shown as a percentage. Each site-day error result is an error sample. Daily error samples are aggregated into cumulative daily error values for each day of the month. The last cumultive dialy error value is the month-end error (MEE).

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 each month. Based on 3,434,166 audit measurements gathered throughout 2019, NetForecast validates with 95% statistical confidence that the typical site was under counted by -4.74%.

References

1. Jonathan Falk, *NERA, Assessment of NetForecast, Inc. Methodology* March 2018.
2. Sevcik, Wetzel, and Lacy, *General ISP Data Usage Meter Specification and Best Practices* NetForecast Report 5119, May 2015.

About the Authors

Andrew Lacy, NetForecast's director of development, has a 30-year track record delivering solutions to complex business and technical requirements. He has extensive experience as a development leader designing, building, and deploying products. He has a strong technical background in data communications, game software, server-based gaming, embedded systems, server software, web, and database design.

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 co-inventor of three patents on application response time prediction and congestion management. He pioneered Internet usage tracking techniques and invented the Apdex methodology.

Rebecca Wetzel is a director at NetForecast, and a data communications industry veteran. She helped realize the commercialization of the Internet in its formative years and worked to design and market some of the Internet's first value-added services. She has spent over two decades as an Internet industry analyst, consultant, writer, and advisor to numerous startups.

Additional information is available at:
www.netforecast.com

NetForecast and the NetForecast logo are registered trademarks of NetForecast, Inc.

APPENDIX A – Yearly Apdex

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. 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 provde 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.

To provide clear insight into accuracy, NetForecast applies the Application 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.

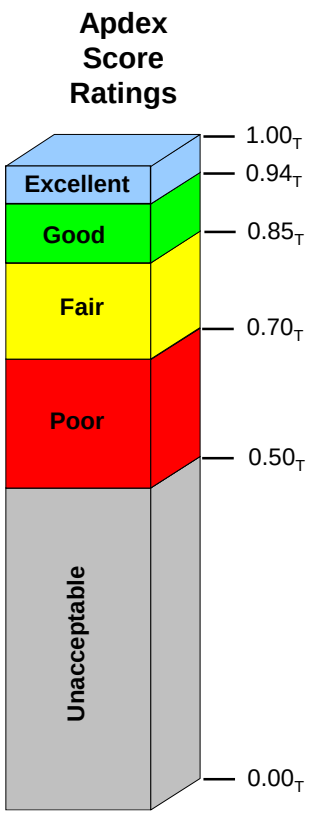

It is a methodology promoted by vendor, enterprises, and professionals who are members of the Apdex Alliance. The Alliance uses the rating scale (left) to communicate scores indicating excellent to unacceptable performance. See www.apdex.org to learn more.

Apdex converts many values into a simple meaningful number that properly reflects the user's perspective of performance achieved relative to a performance target. Assuming a +/-1% specification month-end error, samples are placed into the following categories each month:

In Spec/Compliant: Number of samples within the +/-1% meter specification. These samples clearly meet the goal.

Minor Under Reporting: Number of samples below -1% but greater than -5% (undercounting). Undercounting is outside the specification; however, some modest undercounting can be tolerated since the subscriber is not harmed by some traffic not being counted.

Major Under Reporting: Number of samples below -5% (undercounting). Major undercounting is outside the specification and because it is so far from the actual traffic, even though it benefits the subscriber, it nevertheless undermines the confidence in the accuracy of the meter system.

Over Reporting/Incorrect: Number of samples above +1% (overcounting). Any overcounting 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. The Apdex formula is the numbers of in spec results, plus ¾ of the minor under reporting results, plus 1/2 of the major under reporting results, plus none of the incorrect results, divided by the total number of samples.

$$\text{Apdex[S]} = \frac{\text{In Spec} + (\text{Minor Out of Spec Low} * 0.75) + (\text{Major Out of Spec Low} * 0.5)}{\text{Total samples}}$$

Note: Incorrect samples have zero value