Comcast Internet Usage Meter 2017 Accuracy Audit Report

Peter Sevcik, Andrew Lacy, Rebecca Wetzel
June 2018

Executive Summary

NetForecast has independently and continuously audited the Comcast Internet data usage meter since 2009 [1-6]. This 7th report continues that process with new metrics that include the statistical confidence of the results. Comcast again engaged NetForecast to validate the accuracy of its data usage meter system in 2017. 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 2017, 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.

The Comcast meter met the accuracy specification at month-end for 90.5% of the site months. Out-of-spec low (undercounting) occurred in 8.7% of the site-months, while out-of-spec high (overcounting) occurred in 0.8% of the site months as Figure 1 shows. **NetForecast validates with 95% statistical confidence that in 2017, across all measured sites, out-of-spec monthly errors had an overall impact of undercounting by 4.1%.**

Daily meter errors resulted in an overall Apdex score of 0.95, a rating of "Excellent" for the year (see Appendix A for information on Apdex).

NetForecast gathered usage data at end-user premises throughout the year. The findings in this report are based on 2,868,900 audit measurements which generated 11,954 daily error results. The accuracy assessment in this report is based on the minimum number of samples required for statistical reliability.

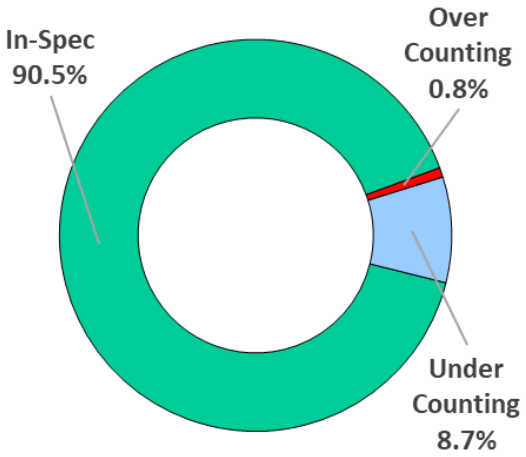

Figure 1 – Comcast Meter Month-End Error Distribution in 2017

Our audit methodology was recently assessed by NERA Economic Consulting [7]. NERA concluded that the measurement technology is unbiased and the analytical process is reliable, with appropriate confidence intervals.

If you are interested in learning why your home data usage is excessive or want to check the Comcast meter yourself, we suggest reading our usage investigation guide [8].

NetForecast Report NFR5128

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.comcast.com. After signing in and selecting "Manage Internet," scroll to the bottom of the page to see a Monthly Data Usage chart as shown in Figure 2. Subscribers can view usage for the current month as well as a rolling 6-month history. Mousing over any month shows the exact usage value in gigabytes (GB).

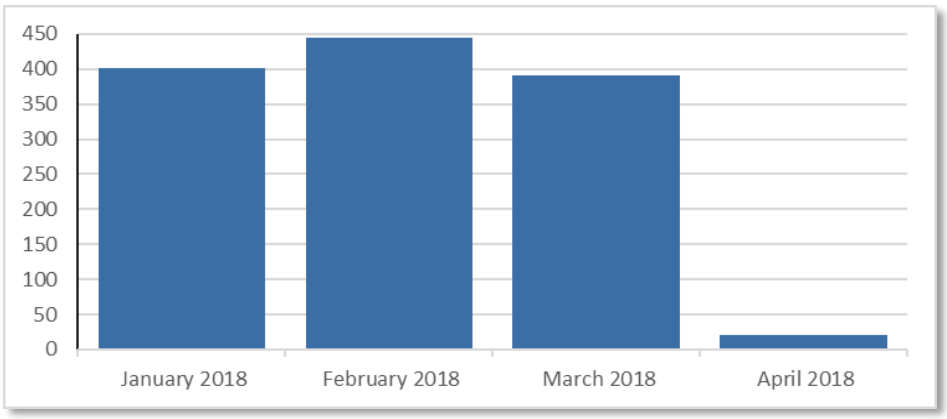

Figure 2 – Cumulative Monthly Usage View

Selecting "Show Table Data" brings up information on the current month and previous 6 months as shown in Figure 3.

Month	Total Monthly Usage	Total Monthly Overage
October 10/01/2017 - 10/31/2017	440 GB	0 GB
November 11/01/2017 - 11/30/2017	389 GB	0 GB
December 12/01/2017 - 12/31/2017	459 GB	0 GB
January 01/01/2018 - 01/31/2018	401 GB	0 GB
February 02/01/2018 - 02/28/2018	444 GB	0 GB
March 03/01/2018 - 03/31/2018	391 GB	0 GB
April 04/01/2018 - 04/30/2018	20 GB	0 GB

Figure 3 – 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 4 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 2 and 3.

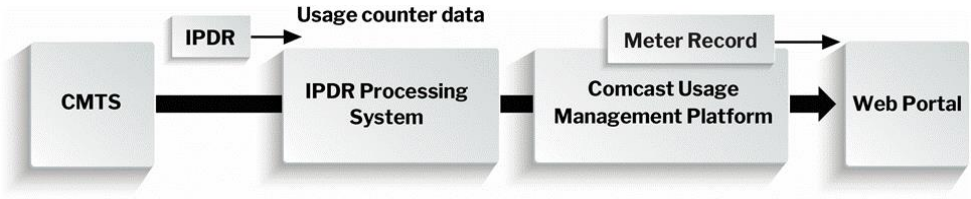

Figure 4 - How the Meter Data Is Processed

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 5.

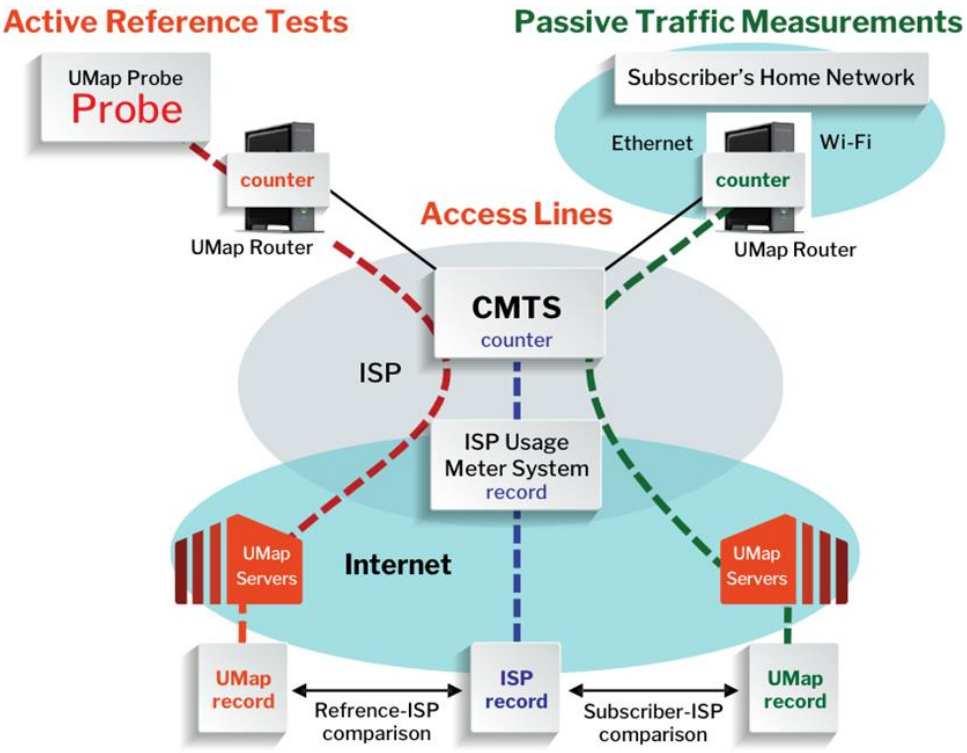

Figure 5 – The NetForecast Methodology

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 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 5. 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 UMap and 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 overcounting. If the error is zero, the meter is as accurate as it can be, and if the error result is negative, the meter is undercounting. 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)..

NetForecast Meter Accuracy Assessment

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 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* [9].

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 a cumulative daily and month-end basis.	NetForecast validates that the meter met the +1% but not the -1% accuracy goal.	+1% Yes -1% No

NetForecast analyzed daily as well as monthly error values. Month-end error analysis reflects the accuracy of the value that appears on a subscriber's monthly bill, while cumulative daily error analysis reflects the accuracy of meter values subscribers can view on Comcast's portal each day.

How well did the meter perform at Month-End?

On average, each site-month has some chance of being out-of-spec high or out-of-spec low as shown in Figure 1. Site months that are within spec as shown in Figure 1 do not apply in this analysis. Within spec means that Comcast met the accuracy goal, thus requiring no further investigation.

NetForecast analyzed the impact of out-of-spec month-end errors both in quantity and magnitude normalized across all sites and all months in the year. The extent of resulting high (overcounting) or low (undercounting) error conditions applied to all site errors shown below in Figure 7.

When looking only at the out-of-spec results across all sites and all months, we can expect with 95 percent confidence that the average error is -4.3% when out-of-spec low (undercounting), and +0.2% when out-of-spec high (overcounting). Therefore, over a year the typical site was undercounted by 4.1%.

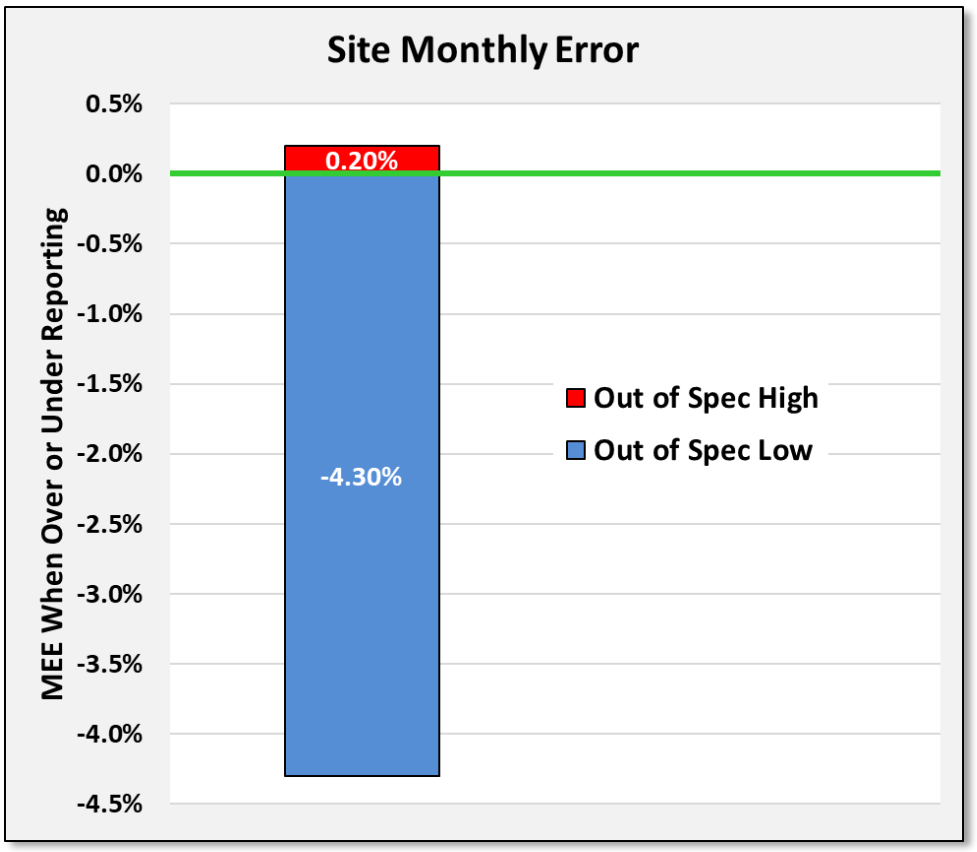

Figure 7 – 2017 Meter Accuracy at Month-End

Meter inaccuracy is typically caused by lost counts or miscounts. Lost counts occur when values are not registered or are dropped during processing, causing a meter system to undercount. Miscounts are counting errors that can cause over or undercounting. Lost counts are much more common than miscounts. For this reason, almost all meter systems—including the Comcast meter system—tend to undercount, which benefits the subscriber.

How well did the meter perform on a cumulative daily basis?

NetForecast applied the Apdex methodology (see Appendix A) to daily error results to determine the meter's overall daily ***accuracy score*** and associated ***accuracy rating***. 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 days were within the specification goal of +/-1%. Site-day errors above +1% negatively impact the scores much more significantly than error below -1%.

NetForecast has measured and documented the cumulative daily error of Comcast's data usage meter system every month since 2012. Figure 8 shows the 6-year Apdex score history. Values below Excellent were overwhelmingly due to undercounting (i.e., count errors under -1%). The 6-year track record using this metric shows a steady improvement of the meter system over time.

During 2017, cumulative daily errors consistently improved, resulting in an annual average Apdex score of 0.95, a rating of Excellent for the year.

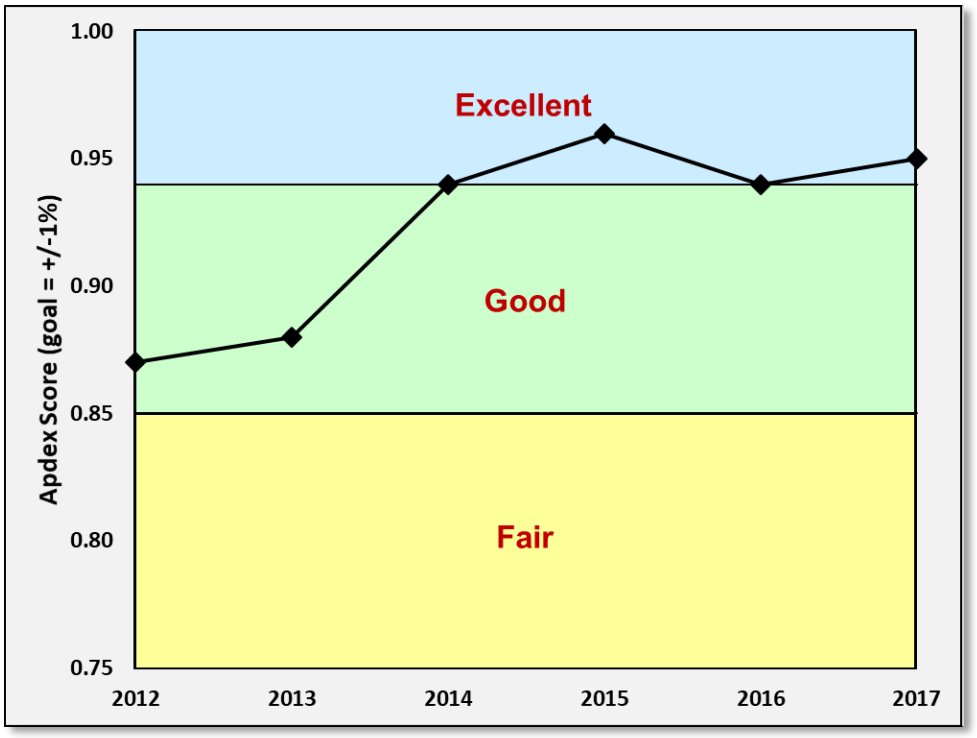

Figure 8 – Daily Meter Apdex Scores Over 6 Years

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. NetForecast validates that the Comcast Internet data usage meter achieves an average daily Apdex accuracy rating of Excellent with 90.5% of the sites within specification on a month-end basis, with a bias toward under-counting.

Our measurement results indicate that subscribers should be able to rely on the accuracy of Comcast's Internet data usage meter.

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, *Fourth Accuracy Assessment of Comcast's Data Usage Meter* NetForecast Report 5120, June 2015.
5. Sevcik, Wetzel, and Lacy, *Comcast 2015 Internet Usage Meter Accuracy Audit Report* NetForecast Report 5121, July 2016.
6. Sevcik, Wetzel, and Lacy, *Comcast 2016 Internet Usage Meter Accuracy Audit Report* NetForecast Report 5126, June 2017.
7. Jonathan Falk, *NERA, Assessment of NetForecast, Inc. Methodology* March 2018.
8. Sevcik and Wetzel, *Understanding and Counting Your Household's Internet Usage* NetForecast Report 5127, May 2018.
9. Sevcik, Wetzel, and Lacy, *General ISP Data Usage Meter Specification and Best Practices* NetForecast Report 5119, May 2015.

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 congestion management. He pioneered Internet usage tracking techniques, and invented the Apdex methodology. He can be reached at peter@netforecast.com.

Andrew Lacy is 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.

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

APPENDIX A – Apdex Applied to Daily Meter Accuracy

Cumulative daily usage view is the sum of usage by day from month start to month end. The typical usage bar 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.

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.

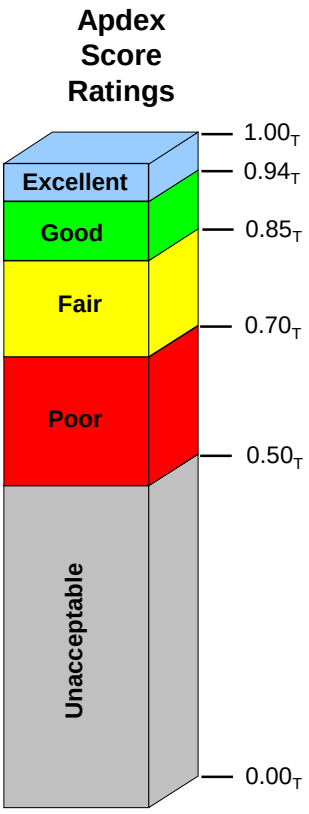

It is a methodology promoted by vendor, enterprises, and professionals who are members of the Apdex Alliance. The Alliance uses the rating scale (left) to communicate scores indicating excellent to unacceptable performance. See www.apdex.org to learn more.

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 cumulative daily 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 below -1% (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.

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 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:

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

Target T or specification is always shown with the score

APPENDIX B – Glossary of Terms

Apdex	A methodology that integrates many measurement samples into a single score representing how well a goal was met. More information is available at www.apdex.org.
Apdex Rating	The Apdex "grade" given to a range of Apdex scores indicating how well the system met the performance objective. Standardized ratings are: Excellent, Good, Fair, Poor and Unacceptable.
Apdex Score	Result of the Apdex formula which is a value between 0 (no samples met the goal) to 1 (all samples met the goal).
Compliant	Apdex count or "bucket" of values that meet the spec.
Cumulative Daily Error	The error each day in a month that is calculated on the sum of ISP values from month start to that day compared to the sum of NetForecast values from month start to that day. The sums increase due to additional subscriber usage over the month.
Gigabyte	2^{30} bytes = 1024 x 1024 x 1024 bytes = 1,073,741,824 bytes
Incorrect	Apdex count or "bucket" of values that are out-of-spec high.
Marginal	Apdex count or "bucket" of values that are out-of-spec low.
MEE	See Month-End Error
Meter Error	Percentage deviation of the ISP's meter count relative to the NetForecast meter count. Positive error indicates the relative amount the ISP overcounted. Negative error indicates the relative amount the ISP undercounted.
Month-End Error	The overall error for all usage counted by the ISP and NetForecast per site. The last cumulative daily error in the month is the month-end error.
Out-of-Spec High	Positive error that is greater than the positive boundary of the meter spec (e.g., +2% error when spec is +/-1%).
Out-of-Spec Low	Negative error that is less than the negative boundary of the meter spec (e.g., -2% error when spec is +/-1%).
Payload	Data carried in a packet on behalf of an application or user.
Within Spec	Error values that are inside the ISP defined acceptable error range (e.g., +/- 1%).