

Comcast Internet Usage Meter 2016 Accuracy Audit Report

Peter Sevcik, Andrew Lacy, Rebecca Wetzel
June 2017

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. NetForecast has published 5 reports on Comcast meter accuracy since 2009 (see references 1 through 5). This, the sixth report, details the accuracy results from measurements at 55 Comcast service locations for the 12-month period from January through December 2016.

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. Using the (+/-) 1% accuracy goal, the Comcast data usage meter's average daily accuracy score was 0.94 on a zero-to-1 scale, with 1 the highest possible score. This results in an overall rating of "excellent" for the twelve-month measurement period. On a month-end basis, the meter did not over-count, although it did experience occasional under-counting, which means that some user traffic did not register on the meter.

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 "Devices," a new page shows a card for each device supported by Comcast. Scroll to the bottom of this page to see the Data Usage Overview displayed as shown in Figure 1. Subscribers can view usage for the current month as well as a rolling 6-month history.




Figure 1 – Cumulative Monthly Usage View

NetForecast Report
NFR5126

©2017
NetForecast, Inc.

Selecting “Show Table Data” brings up information on the current month as previous 6 months as shown in Figure 2.

Month	Total Monthly Usage	Total Monthly Overage
December 12/01/2016 - 12/31/2016	441 GB	0 GB
January 01/01/2017 - 01/31/2017	480 GB	0 GB
February 02/01/2017 - 02/28/2017	428 GB	0 GB
March 03/01/2017 - 03/31/2017	444 GB	0 GB
April 04/01/2017 - 04/30/2017	566 GB	0 GB
May 05/01/2017 - 05/31/2017	497 GB	0 GB
June 06/01/2017 - 06/30/2017	253 GB	0 GB

Figure 2 – Month-End Usage Table 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 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 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 single usage meter that aggregates the usage from all the devices on the account, or a separate meter for each device (depending on region; Comcast is currently rolling out an account-level aggregated view).

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 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 UMapSM 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. Once installed, the routers register with the NetForecast UMap management and reporting system.

The UMap system operating on the Comcast network performs two types of measurements: active and passive as shown in Figure 4. This report is based on results from 5 active and 50 passive measurement sites.




Figure 4 – 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 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 the Comcast meter system 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 meter daily traffic measurement records.

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

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.

NetForecast Meter Accuracy Assessment

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. 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 [6]

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 accuracy goal with an overall Excellent rating. Details of the assessment are described below.

This section presents NetForecast’s findings regarding Comcast’s accuracy performance history relative to its error goals. Details of the analysis methodology are described in Appendix A. The accuracy assessment presented here assumes that the Comcast meter operates uniformly across Comcast’s entire service area.

The data analysis answered the following questions:

- How well did the meter perform overall?
- How often was the meter in spec?
- How far was the meter out of spec?

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 overall?

NetForecast applied Apdex methodology (see Appendix A) to monthly as well as daily error results to determine the meter’s overall **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”.

The Apdex method groups error values into three buckets: compliant (meets the spec), marginal (out-of-spec low) and incorrect (out-of-spec high). The number of samples in each bucket is then converted into a ratio that ranges from zero (all values are out of spec because they are too high) to 1 (all values are within spec). The zero-to-one value is the score (i.e., numeric grade). Using the Apdex methodology, scores correspond to accuracy ratings of Excellent, Good, Fair, Poor, or Unacceptable.

Figure 5 shows accuracy scores based on daily error values for all of 2016. The average accuracy score for the year was 0.94, which rates as “Excellent”. The Comcast meter operated in the Excellent range most of the year, then slipped into the “Good” range at the end of the year.




Figure 5 – Daily Meter Accuracy History

How often was the meter in spec?

Figure 6 shows the frequency with which sites were within and outside of the +/-1% accuracy goal each month. Under-counting occurred during 9 out of the 12 months, and there were no instances of over-counting. The number of sites within specification at month end was 92% across the year.




Figure 6 – How Often Sites Were Within/Out of Spec

How far was the meter out of spec?

It is highly unusual for an Internet data usage meter to be completely accurate all of the time. When meter results do diverge from the specification, it is important to determine the average error values above and/or below the specification to understand the magnitude of the error.

Figure 7 shows the average magnitude of meter errors during the measurement period. Each colored square illustrates the error associated with the instance of the type of errors shown in Figure 6.

Under-counting, which occurred in 9 of the months, averaged about -1.9% over the test period. Over-counting did not occur in 2016.




Figure 7 – How Much Sites Were Within/Out of Spec

It is useful to put out-of-spec errors into perspective. The overall effect of errors depends on the product of the number of times (Figure 6) the condition occurred and the magnitude (Figure 7) of the condition. Under-counting (negative) errors benefit a subscriber. Over-counting (positive) errors can adversely affect a subscriber who exceeds their data plan, however, since no over-counting was detected, it is highly unlikely that any Comcast subscribers were adversely impacted in 2016 due to meter inaccuracy.

Meter inaccuracy is caused by lost counts or miscounts. Lost counts, which occur when values are not registered or are dropped during processing due to resource limitations, cause under-counting. Miscounts are counting errors that can cause over or under-counting. Lost counts are much more common than miscounts. For this reason, almost all meter systems—including the Comcast meter system—exhibit a bias towards under-counting that benefits the subscriber.

Useful Information if You Want to Do Your Own Counting

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

Avoiding binary versus decimal math confusion

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

Binary			Decimal	
KB	Kilobyte	1,024	Thousand (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 one wrongly to conclude that the ISP meter is overcounting.

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

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

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

Details that may seem minor can mess up your counts. For example, we recently discovered a home router that appeared to count properly, but it only counted usage for devices in the DHCP table at start up. Usage by devices added to the network after the router booted went uncounted. Rebooting the router brought the new devices into the counts from the reboot onward. The subscriber reasonably concluded that the ISP was overcounting, but in fact, months had passed since the last router reboot, and new devices were introduced into the home during that period. These new devices generated significant usage that the router did not count but the ISP did. Properly measuring home usage requires technical know-how, careful attention to process, and patience.

Understanding Unexpected Traffic

If you look closely at your Internet usage, chances are you will see unexpected bytes register on your meter. Most Internet users do not 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 revealed unexpected high consumption caused by cloud storage and file sharing services such as Google Drive, OneDrive, Box, Dropbox, and iCloud. Some cloud services can enter an endless synchronization loop, which dramatically increases traffic counts. We have observed this when very large files are shared across multiple users. If you use one of these cloud services and experience 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 only to PCs, but also 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. For convenience, these automated connections are invisible to the user, so you may be unaware of the traffic generated. In addition, a neighbor or “wardriver” may use your connection without your permission or knowledge.

Security related issues can also contribute to unexpected traffic. A PC could be hijacked and generate traffic that has nothing to do with any user in your home. Specifically, botnets and malware that have infected one or more computers in the home network can be leveraged for outbound Denial of Service attacks against targets on the Internet and/or can be used as proxies to route traffic for bad actors. Bad actors can also use UDP-based amplification attacks to exploit vulnerabilities on a home router, resulting in high usage consumption [7].

Excessive up usage during times when no known backups were occurring is an indication that your home is involved in a UDP-based amplification attack. If your ISP shows usage

by day or hour separately for down and up traffic, you can check for anomalously large up usage.

Lastly, usage related to inbound probes and attacks, which normally are filtered by the home router firewall, are counted in the ISP usage meter.

The most likely source of unexpected traffic, however, is PC software and mobile applications. Most popular software has automated update features that 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 are not 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 pushed overnight to a smart phone or tablet for viewing or listening the following morning.

Assume each night's download 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 streaming service—and you may get more bytes than you expected 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 [8].

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 sites from January through December 2016, NetForecast validates that the Comcast Internet data usage meter achieves an average daily accuracy rating of Excellent, and 92% of the sites were within specification on a month-end basis, with a bias toward under-counting.

Based on our measurement results, subscribers should be able to rely on the meter's accuracy, however, subscribers should be aware that the meter has shown a slight tendency to record less usage than actually occurred.

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 Internet Usage Meter 2015 Accuracy Audit Report](#), NetForecast Report 5121, July 2016
6. Sevcik, Wetzel, and Lacy, [General ISP Data Usage Meter Specification and Best Practices](#), NetForecast Report 5119, May 2015.
7. US Computer Emergency Readiness Team (US-CERT), [UDP Amplification Attack](#)
8. 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 Authors

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 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.

APPENDIX A – How NetForecast Reports Meter Accuracy

The following is a general high-level description of the NetForecast meter accuracy reporting methodology. Data shown in this appendix does not represent any specific ISP's accuracy performance. Descriptions and examples in this appendix are based on a +/-1% error specification.

Daily and Monthly Error Calculations

The ISP meter portal view shows two views of usage to subscribers:

Cumulative Daily Usage View: 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. 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. 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 the ISP's meter system for the same traffic, to determine whether the comparisons fall within the ISP's target accuracy range. Meter error is therefore calculated on both a cumulative daily basis and a month-end basis.

Given differences in the nature and sample size of the two ISP 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*.

The two views are analyzed in three levels of detail as shown in the table below. Results with an “X” are documented in this report.

	Assessment Method	Cumulative Daily Error	Month-End Error
High-Level Assessment	Apdex scores of error values and accuracy ratings	X	
Goal Achievement: How often the meter meets the specification	Percent of sites in group (Spec: In, Out Low, Out High)		X
Meter Error: How far the meter is out of specification	Average error value in each group		X

High-Level Daily Assessment


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. 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 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 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. Assuming a +/-1% specification 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). Under counting 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% (any 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 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:

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

↑
Target T or specification is always shown with the score

↓
Incorrect samples have zero value

The chart below shows a typical cumulative daily meter error distribution or frequency plot. In the example below, the samples highlighted in green are within a +/-1% specification, while samples in yellow are out-of-spec low (undercounting) and those in red are out-of-spec high (overcounting). Note that the horizontal scale denoting the distribution buckets changes from a bucket every 0.1%, when within spec, to a bucket every 1.0% for out-of-spec values. There is clear bias towards undercounting: more values in yellow vs. red in this example.



Adding the bucket percentages (blue bars) in the green range shows that 88% of the samples are within spec. The sum for bars in the yellow range is 10%, while the bars in the red range add up to 2%, accounting for 100% of the samples. Applying the above Apdex formula to the number of samples shown in this distribution chart yields: $Apdex = 0.88 + 0.10/2 = 0.93$. The distribution has an Apdex score of 0.93 [+/-1%], which has a rating of Good.

Note that the Apdex score must always show the goal or objective in brackets []. The score is 0.93 [based on a goal of +/- 1%].

Achieving the Monthly Goal – How Often the Meter Is Within the Specification

The frequency with which sites fall within or outside the +/-1% accuracy goal each month is determined on a month-end basis. The frequency is described as the percentage of sites within the specification, out-of-spec low, and out-of-spec high. The values are displayed on a color-coded bar chart showing:

- Out-of-spec Overreporting – Red
- Within spec – Green
- Out-of-spec Underreporting – Yellow.

Monthly Meter Error – How Far the Meter Is Outside the Specification

The final view shows the average magnitude of the month-end meter errors within each error group during the measurement period. An error value is associated with each bar chart color in the “How Much Sites Were Within/Out of Spec” chart.

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.
DHCP	Dynamic Host Configuration Protocol
Error Group	Sites grouped as follows: Within Spec, Out-of-Spec Low, and Out-of-Spec High.
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
Meets the Spec	Error values that are inside the ISP defined acceptable error range (e.g., +/- 1%).
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.
Reporting Period	The time period of an assessment report. Typically, this is a year for annual accuracy reports.
TCP	Transmission Control Protocol used in most devices connected to the Internet.
UDP	User Datagram Protocol
Within Spec	Error values that are inside the ISP defined acceptable error range (e.g., +/- 1%).