Cox Internet Usage Meter 2018 Accuracy Audit Report

Peter Sevcik, Andrew Lacy, and Rebecca Wetzel
May 2019

Executive Summary

NetForecast began measuring the accuracy of the Cox Internet data usage in 2010 and has continuously audited the accuracy of meter since 2016. This report covers the 12-month period from January through December 2018. The Cox meter accuracy specification states that the meter should correctly measure traffic within plus-or-minus (+/-) 1% on a month-end and daily basis. Significant undercounting occurred throughout 2018, which benefited Cox's subscribers. There were infrequent and small overcounting occurrences, however, there was no observed overcounting above the +1% specification over the course of a full month.

For Cox, NetForecast performs both passive measurements of real user traffic in subscriber homes, and active reference tests in which a NetForecast probe 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 Cox for each location, and compares NetForecast's measurements with Cox's records.

NetForecast gathers usage data throughout the year. The accuracy assessment in this report is based on 10,800,012 audit measurements, exceeding the minimum number of samples required for statistical reliability. 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 2018, across all measured sites, out-of-spec monthly errors had an overall impact of undercounting by 7.73% as Figure 1 shows. This undercounting benefited Cox subscribers.

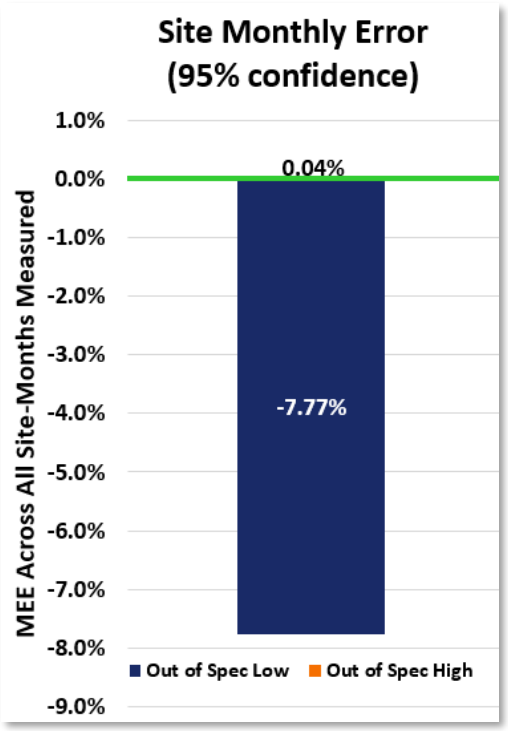

Figure 1 – Site Monthly Error in 2018

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

Month-end meter errors resulted in an overall Apdex score of 0.89, a rating of "Good" for the year (see Appendix A for information on Apdex).

NetForecast's audit methodology was recently 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 NFR5130

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/under-counting is reflected in the monthly error calculations, which quantify how much the meter would be affected (with 95% confidence) when a site over or under reported. Only out-of-spec site months (those with values greater than +1.0% and less than -1.0%) are used in this metric.

During 2018, the Cox meter met the accuracy goal at month end 65% of the time as Figure 2 shows. Given the margin of error, the sample size, and the standard deviation, the Cox meter could be expected to over count 0.2% of the time and under count 34.8% of the time in 2018.

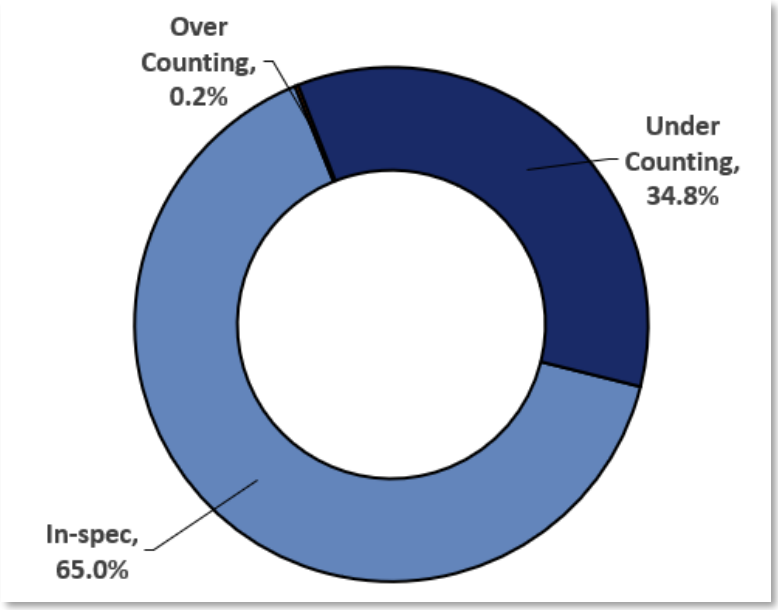

Figure 2 – Cox Meter Month-end Error Frequency Distribution in 2018

As Figure 1 shows, when a site under counted, the expected underreporting amount was -7.77%, and when a site over counted, the expected overreporting amount was +0.04%.

Under counting 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. Over counting is rare compared to under counting 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 Cox's data usage meter system for every measurement site in every month in 2018. Figure 3 shows the Apdex score history for all years the Cox meter has been continuously measured with the current NetForecast methodology. Figure 8 shows the Apdex score history. Values below Excellent are overwhelmingly due to undercounting (i.e., count errors under -1%).

Although exhibiting a downward trend in 2018, the month-end errors resulted in a yearly Apdex score of 0.89, a rating of Good for the year.

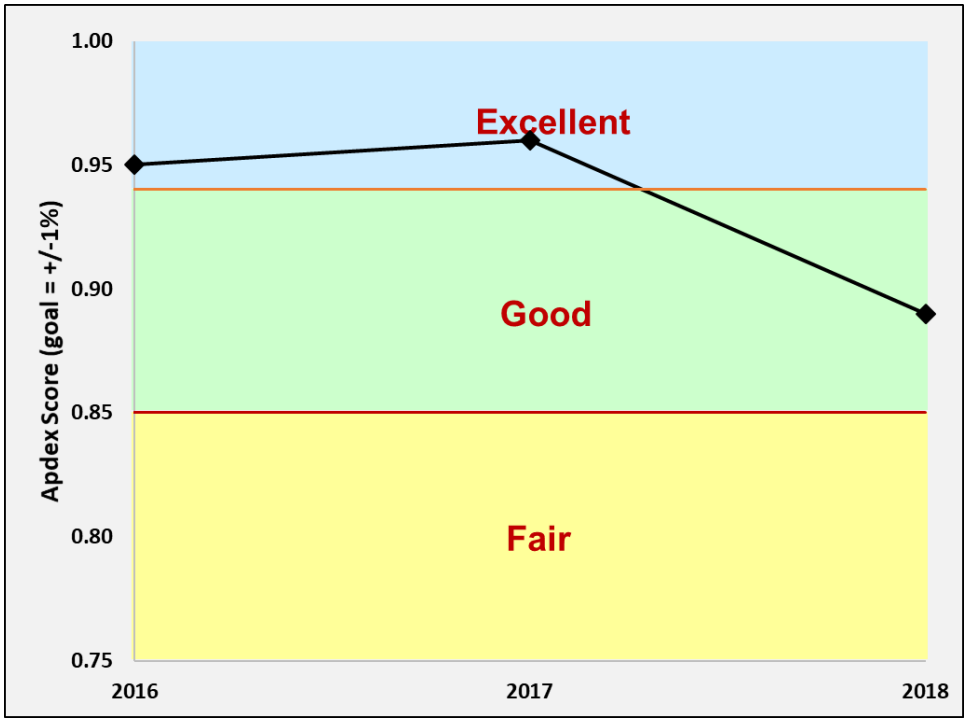

Figure 3 – 2016 through 2018 Yearly Apdex Scores

Meter Accuracy Factors

Cox worked with NetForecast to create a Cox-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 Cox'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	Cox 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

The Cox Usage Meter

The Cox Internet data usage meter provides subscribers with information about how much traffic has crossed their residential Internet connection. The meter sums traffic to and from the Internet and displays a data consumption summary for the current billing period as shown here.

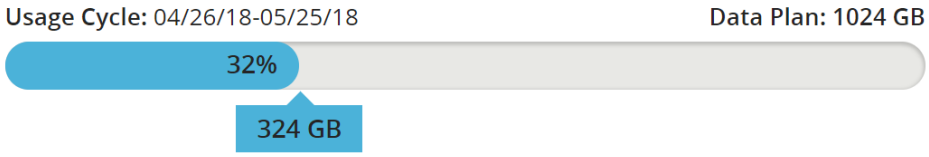

Figure 4 – Summary Usage View

Cox's Internet usage reports are accessible via the Cox customer portal at http://www.cox.com/. If you are a subscriber, you can view your Internet usage information by logging in at the sign-in prompt at the top right of the portal menu bar. This opens the "My Account" window. The data usage meter link is in a box labeled "My Internet". Click on "View More" within the My Internet box to reveal the data usage meter icon.

Scrolling down the usage meter page shows usage over time, with selections for *monthly*, *current daily*, and *past daily* views shown below. All counts are rounded to the nearest GB.

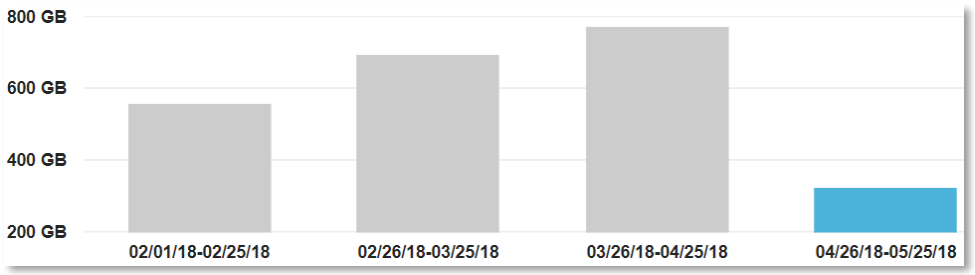

Figure 5 – Cumulative Monthly Usage View

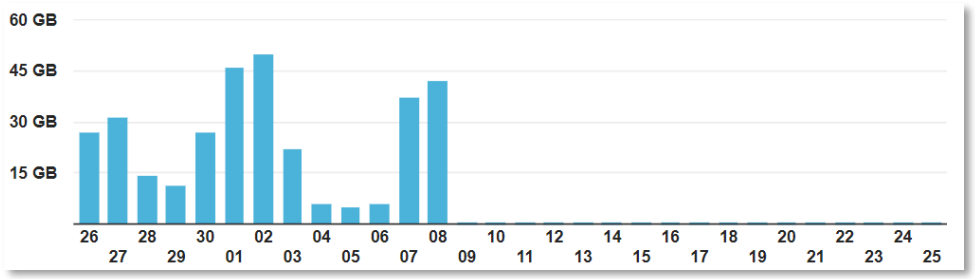

Figure 6 – Current Daily Usage View

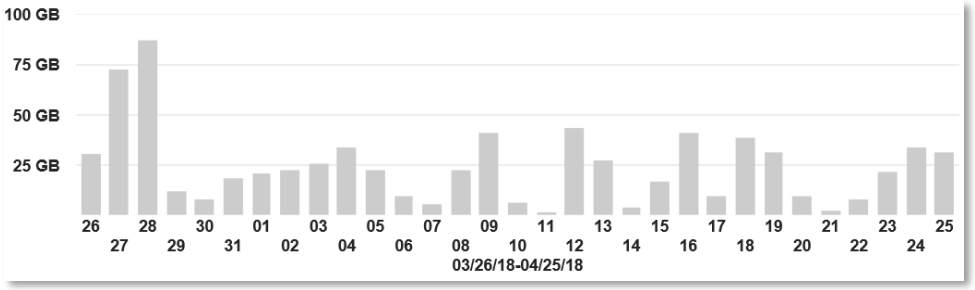

Figure 7 – Past Daily Usage View

How the Cox Data Usage Meter Works

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

The CMTS counts downstream and upstream traffic for each subscriber cable modem it serves. Downstream traffic flows from the Internet to the subscriber, and upstream traffic flows from the subscriber to the Internet. The measurement sites were supported by Cisco uBR10000 series CMTSs. Each CMTS periodically reports the down and upstream counts in an Internet Protocol Detail Record (IPDR) as shown in Figure 8.

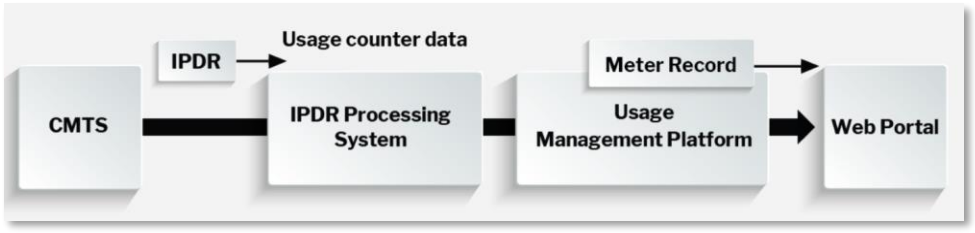

Figure 8 - How the Meter Data Is Processed

From the CMTS, IPDRs are sent to the IPDR collector, which collects, processes, and stores IPDR data streamed from CMTSs. IPDR processing involves time normalization and usage data mediation, and the IPDR Collector cross checks the completeness and correctness of IPDR data to detect anomalies that can result in inaccurate Internet usage data.

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

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

What the Meter Shows

There are a numbers steps between the time a subscriber's packet moves through the cable modem and the counts appear on the Cox subscriber portal. CMTS traffic counts are generally reported every 15 minutes. The IPDR processing system aggregates the counts and summarizes them by hour. The usage management platform database receives the updates and prepares the results to be shown on an hourly basis. The meter on the portal generally updates within three hours after traffic passes over the network.

NetForecast Meter User Experience Assessment

Goals were also set for the user's experience of the data usage meter as displayed on the Cox subscriber portal. The following table summarizes the goals for the user experience specification factors, and Cox's performance relative to those goals.

Meter Factor	Cox Goal	NetForecast Assessment	Goal Met
User Experience Factors			
Timeliness	Usage occurring in an hour will update onto the web portal by noon the following day.	The meter updates by 7AM 90% of the time, and always by noon.	Yes
Granularity	Usage shown in whole GB (single number).	The usage meter displays whole GB values.	Yes
Mathematical Consistency	Portal details add up to summary usage value within +/- 2% of the total value.	The manually tabulated sum of the daily usage values may differ from the monthly value shown on the portal because the daily data is rounded multiple times. The difference is, however, within +/-2% of the total value.	Yes
Accessibility	Usage meter is no more than two clicks after logging into portal.	The usage meter is one click from portal login.	Yes
Availability	Usage meter is available 99% of the time.	This was not assessed; however, the portal was available when accessed.	Yes
Clarity	Comprehensive FAQ pages.	The FAQ pages contain limited, but not comprehensive information about the meter.	No

For a 30-day period, NetForecast recorded the hour during which the portal updated. As Figure 9 shows, 90 percent of the time the Cox meter values updated by 7AM the following day, and the meter always updated before noon, thus exceeding the timeliness goal.

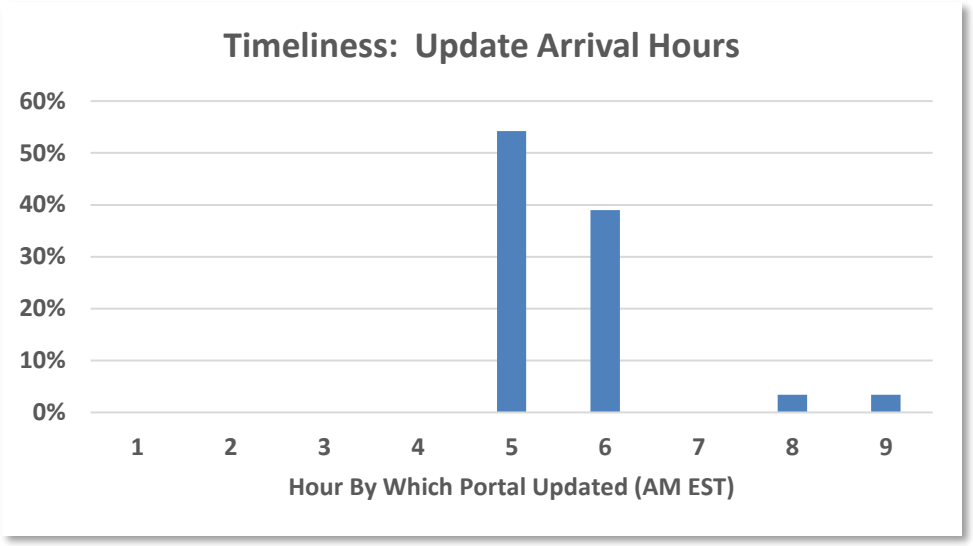

Figure 9 – Meter Update Performance

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 Cox network performs two types of measurements, active and passive, as shown in Figure 10.

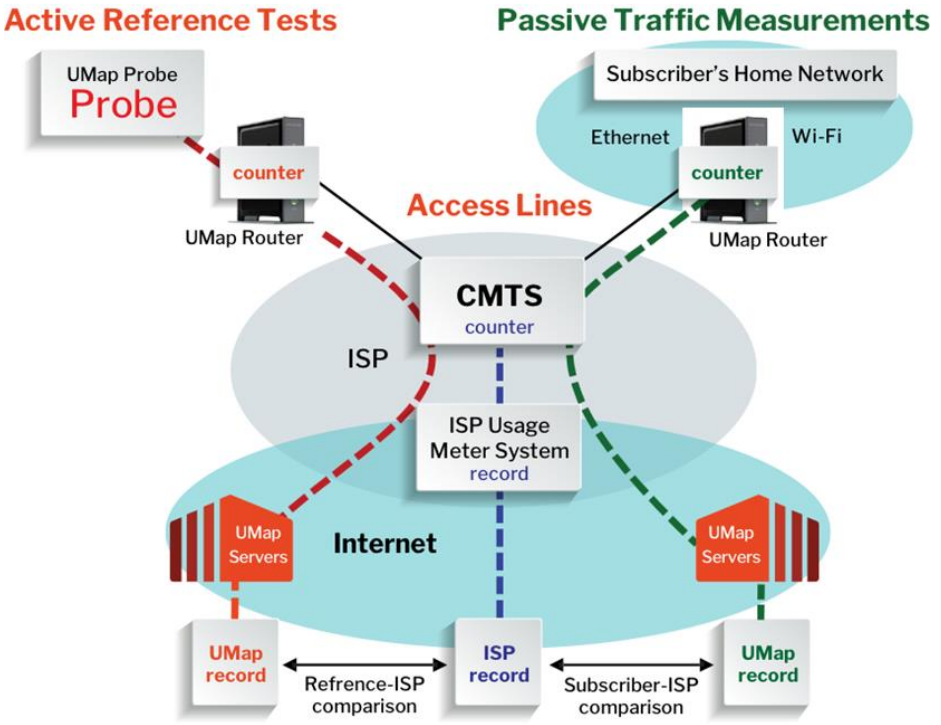

Figure 10 – The NetForecast Methodology

Active Reference Testing

This description applies to the red dashed line of data flow in Figure 10. 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 10. 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 Cox meter system so the same hours are compared. Daily sums are generated for each site. NetForecast applies the following formula to the UMap and Cox usage meter daily traffic measurement records:

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

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).

Conclusions

Cox 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 10,800,012 audit measurements gathered at end-user premises throughout 2018, NetForecast validates with 95% statistical confidence that the typical site was under counted by -7.73%.

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

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 Analytics and has a 30-year track record delivering solutions to complex business and technical requirements. He has extensive experience architecting, building, and analyzing data for Enterprise Database systems. He is also a development leader designing, building, and deploying many different 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 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

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 occasionally to 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 portion of the sample size, the meter cannot be considered accurate.

To provide clear insight, 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.

The methodology is promoted by vendors, enterprises, and professionals who are members of the Apdex Alliance. The Alliance uses the rating scale 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 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 number of in-spec results, plus ¾ of the minor under reporting results, plus ½ 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