



Suddenlink Internet Usage Meter 2015 Accuracy Audit Report

Peter Sevcik, Andrew Lacy, Rebecca Wetzel
April 2016

The Suddenlink data usage meter system measures and reports how much Internet data traffic a subscriber consumes and generates. Suddenlink has engaged NetForecast to independently and continuously audit the accuracy of its data usage meter system. NetForecast published an initial audit of Suddenlink's data usage meter system in April 2013 [1]. This report details the accuracy results from measurements at 60 Suddenlink service locations from January through December 2015.

For Suddenlink, NetForecast performs passive measurements of real-user traffic in subscriber homes, obtains daily usage meter records from Suddenlink for each location, and compares NetForecast's measurements with Suddenlink's records. Results from the 60 subscriber locations are aggregated in this report. NetForecast analyzed accuracy from a month-end as well as a daily perspective.

Suddenlink established a goal for its Internet data usage meter to correctly measure traffic passing through a subscriber's cable modem within plus or minus (+/-) 2% accuracy. Using the (+/-) 2% accuracy goal, the Suddenlink data usage meter's average month-end accuracy score was 0.93 on a 0-to-1 scale, with 1 the highest possible score. This results in an overall rating of "Good" for the one-year measurement period.

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

The Suddenlink Usage Meter

Suddenlink's Internet usage meter provides subscribers with information about how much traffic has traversed their residential Internet connection. Subscribers can view a usage summary of the current billing period (Figure 1), by month (Figure 2), and by day (Figure 3). The reports sum traffic passing from the Internet into the household with traffic passing from the household to the Internet into a single count.

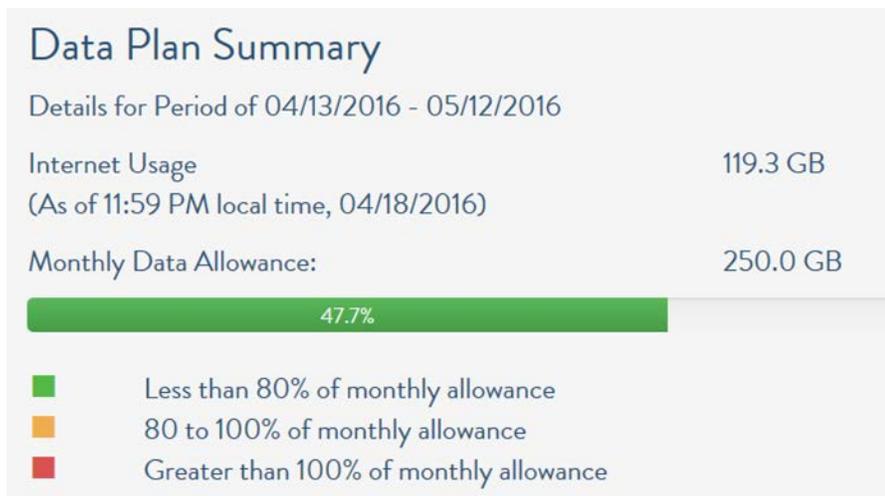


Figure 1 – Suddenlink Cumulative Daily Usage View

Usage reports are available at Suddenlink’s customer portal: <http://www.suddenlink.net/>. Subscribers who have set up an online account at suddenlink.net can access their individual Internet usage report by logging in, selecting *My Account*, then *Online Tools*, and then click *View* under the *Data Plan Usage* section to obtain the summary view shown in Figure 1. Monthly and daily views shown in Figures 2 and 3 are available from a clickable list on the left of each view.

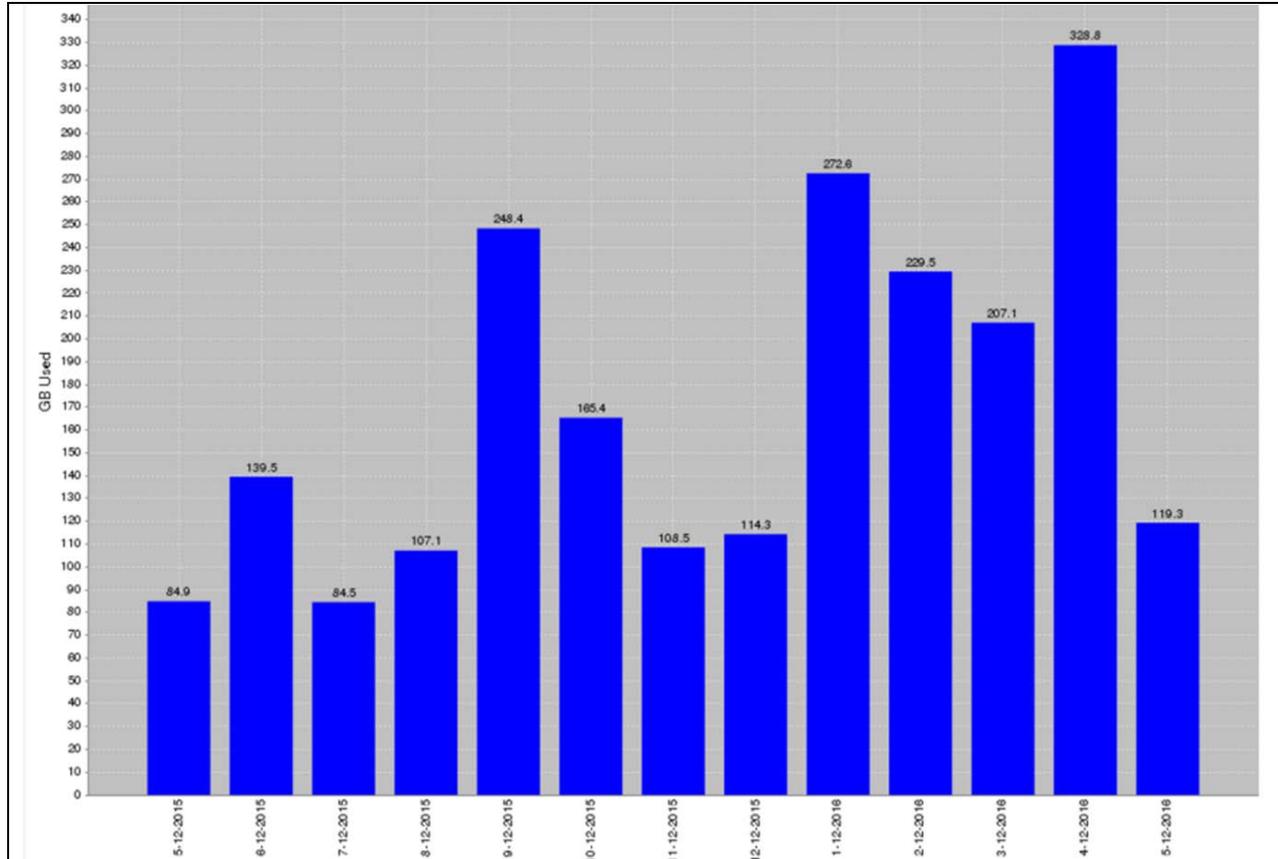


Figure 2 – Month-End Usage View

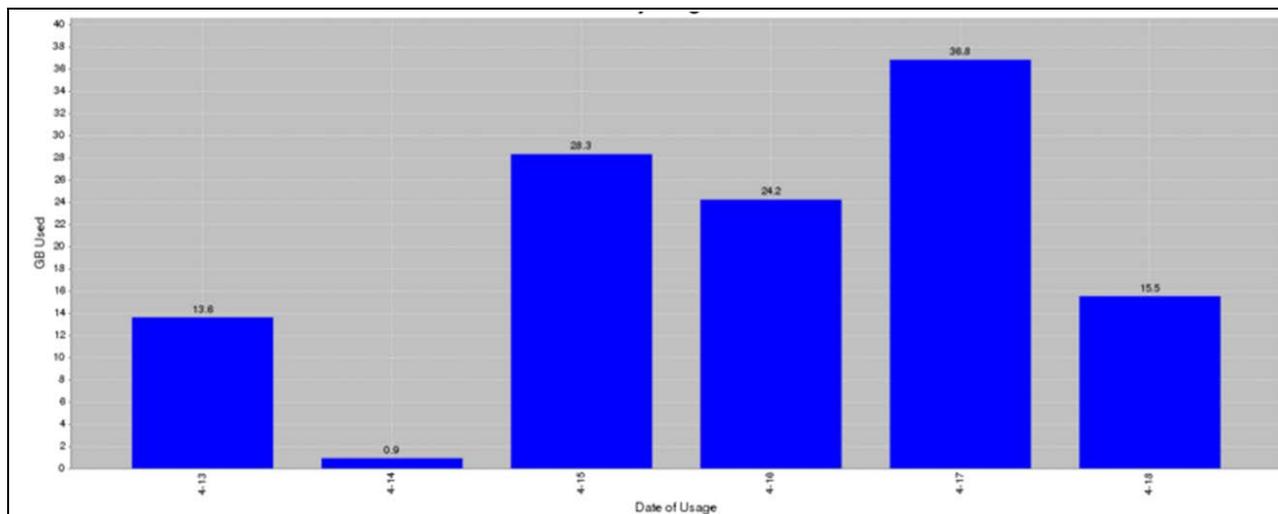


Figure 3 – Daily Usage View

Suddenlink subscribers connect to the Internet through a cable modem at the subscriber's location, 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 Suddenlink'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 down and upstream counts in an Internet Protocol Detail Record (IPDR).

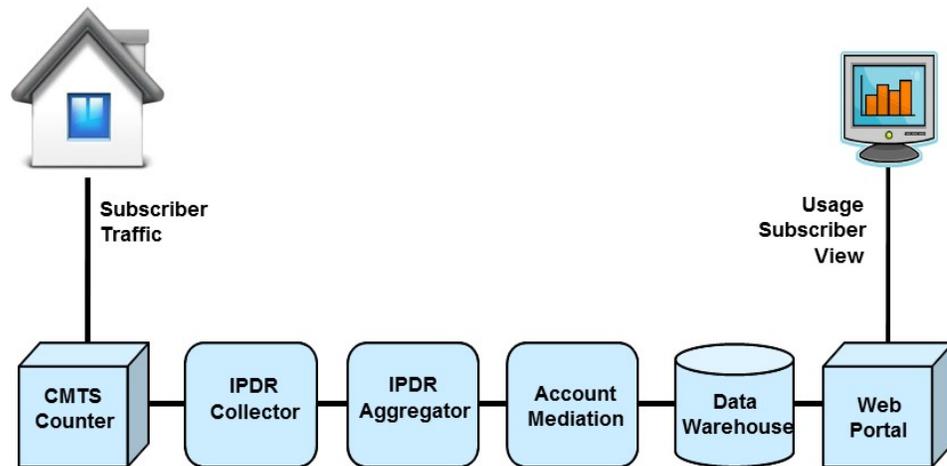


Figure 4 – Usage Meter Process Flow

From the CMTS, IPDRs are sent to an IPDR collection system which performs data remediation to ensure that each IPDR format is correct and that the CMTS counters are incrementing for each cable modem. This prevents faulty data from entering into the system and causing usage reporting issues.

At the next stage, IPDRs from several collectors are aggregated and incremental traffic counts are converted into traffic data. The mediation engine software then associates traffic data with subscriber accounts, and calculates the meter value for each day. Meter data is passed to the data warehouse, which feeds it to the Suddenlink subscriber portal to display the usage views shown in the previous section.

The requirements for how a cable modem communicates with the CMTS and for how subscriber traffic is transported is 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 counted in the IPDR.

NetForecast Internet Usage Meter Accuracy Validation Methodology

All measurements were performed using the NetForecast UMap service delivery platform shown in Figure 5. The UMap system is enabled by measurement and reporting software embedded into customized, fully-featured, wireless home routers supplied and supported by NetForecast. Once installed, the routers register with the NetForecast UMap management and reporting system.

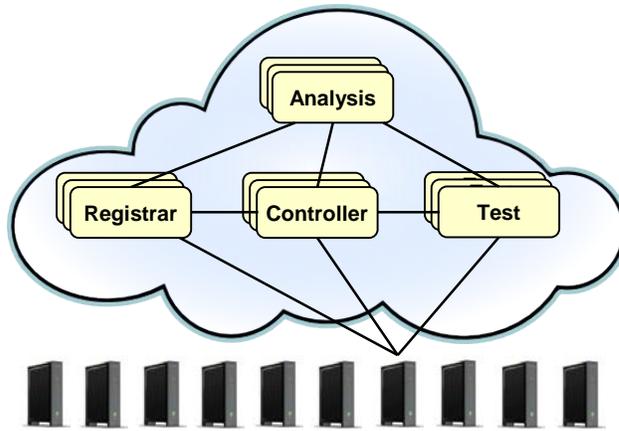


Figure 5 – NetForecast UMap Service Delivery Platform

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

The UMap system operating on the Suddenlink network currently performs passive testing, as shown in Figure 6. This report is based on results from 60 passive measurement sites.

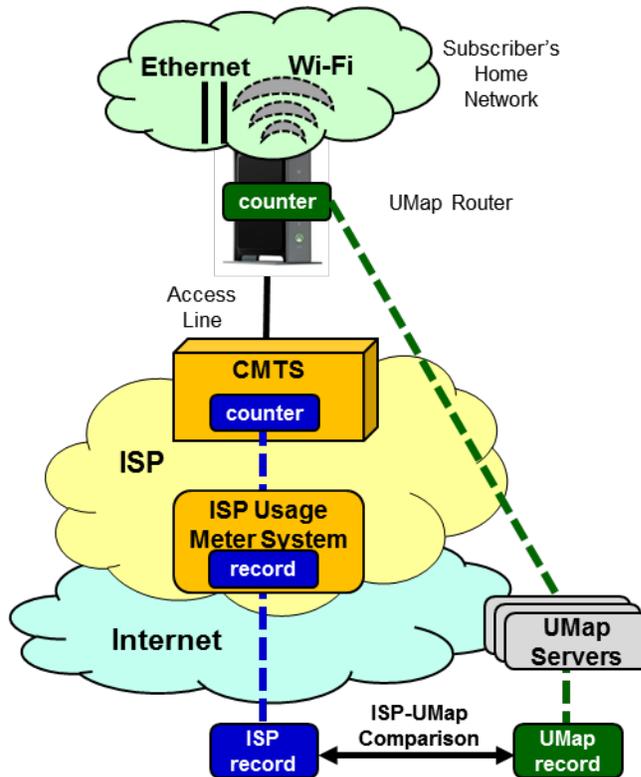


Figure 6 – The NetForecast Methodology

The UMap system continuously measures real-user traffic traversing the home Internet connections of actual subscribers. Measurements take place under real-world conditions—i.e., the measurement system used for Suddenlink accuracy validation is

passive, relying exclusively on the subscriber’s home traffic. This enables measurements to be performed at many locations, thus providing comprehensive geographical coverage. The measurement system provides data for every hour during the month to give comprehensive 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 Suddenlink 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 Suddenlink usage meter daily traffic measurement records.

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

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.

NetForecast Meter Accuracy Assessment

Suddenlink worked with NetForecast to create a Suddenlink-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 Suddenlink’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 [2]

| Meter Factor | Suddenlink Goal | NetForecast Assessment |
|--------------------------|--|--|
| 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 daily basis (local time). | NetForecast validates that the meter records are processed and stored on a daily 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 +/-2% on a cumulative daily and month-end basis. | NetForecast validates that the meter met the accuracy goal with an overall Good rating. Details of the assessment are described below. |

This section presents NetForecast’s findings regarding Suddenlink’s accuracy performance history relative to its error goals. NetForecast began by identifying sites that fell within the middle 90% of sites based on month-end error, and then analyzed data from all of these sites. Details of the analysis methodology are described in Appendix A. The accuracy assessment presented here assumes that the Suddenlink meter operates uniformly across Suddenlink’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 Suddenlink’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 7 shows accuracy scores based on month-end error values for all of 2015. The average accuracy score for the period was 0.93, which rates as “Good”. January - June, the accuracy scores for the Suddenlink meter remained mostly in the “Fair” to “Good” range, but scores improved in the latter half of the year, remaining mostly in the “Excellent” range.

The Suddenlink accuracy score for daily error values across the six months was 0.91, which is also “Good”.

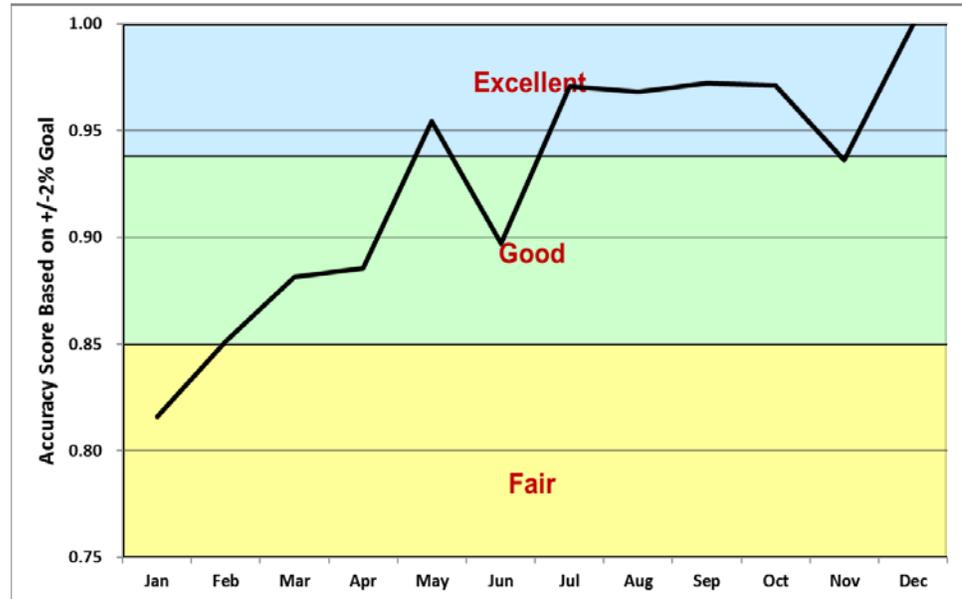


Figure 7 – Month-End Meter Accuracy History

How often was the meter in spec?

Figure 8 shows the frequency with which sites were within and outside of the +/-2% accuracy goal each month. Under-counting occurred every month during the measurement period except December, while over-counting occurred during 5 months, during January, February, March, September, and November.

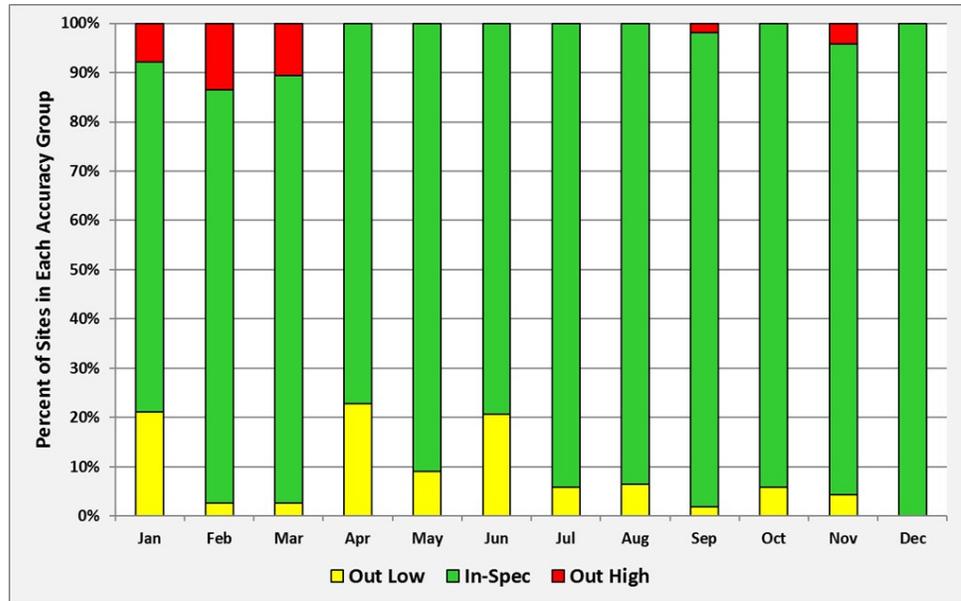


Figure 8 – How Often Sites Were Within/Out of Spec

A summary of the months in which the meter was significantly out of spec is shown in the table below.

| | Metric | Percent of Sites In/Out of Spec | Month |
|-------------------------|-----------------------|---------------------------------|----------|
| Out-of-Spec High | Worst month | 14% | February |
| Within Spec | Average of all months | 88% | NA |
| Out-of-Spec Low | Worst Month | 23% | April |

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 9 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 8.

Under-counting, which occurred consistently almost every month, averaged about -3.6% during 2015. Over-counting, which occurred during 5 months, averaged +2.7% over the year. Over-counting is especially important to track.

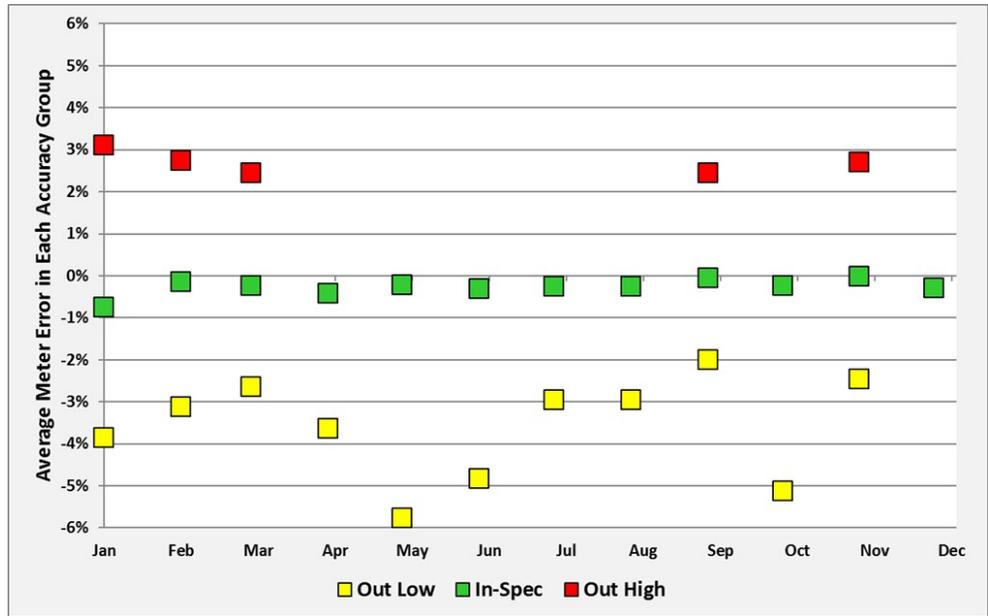


Figure 9 – How Much Sites Were Within/Out of Spec

A summary of the extent by which the meter was in or out of the spec is shown in the table below.

| | Metric | Average Error Value | Month |
|-------------------------|-----------------------|---------------------|---------|
| Out-of-Spec High | Worst month | 3.1% | January |
| Within Spec | Average of all months | -0.3% | NA |
| Out-of-Spec Low | Worst Month | -5.8% | May |

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 8) the condition occurred and the magnitude (Figure 9) of the condition. Under-counting (negative) errors benefit a subscriber. Over-counting (positive) errors can adversely affect a subscriber; however, subscribers who do not exceed their data plan usage allowance are unaffected.

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 over-counting.

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 under-counting. 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 over-counting, 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 [3].

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 3.

Conclusions

Suddenlink established an accuracy goal for its Internet data usage meter to correctly measure traffic passing through a subscriber's cable modem within +/-2% daily and over a month. Based on hourly measurements at 60 sites from January through December 2015, NetForecast validates that the Suddenlink Internet data usage daily and meter month-end accuracy achieves an average accuracy score of Good, although under and over-counting did occur.

Based on our measurement results, subscribers should be able to rely on the meter's accuracy, especially if the accuracy improvement experienced in the latter half of the year continues.

References

1. Sevcik, [Accuracy Assessment of Suddenlink Communications' Internet Usage Meter](#), NetForecast Report NFR5115, April 2013.
2. Sevcik, Wetzel, and Lacy, [General ISP Data Usage Meter Specification and Best Practices](#), NetForecast Report 5119, May 2015.
3. US Computer Emergency Readiness Team (US-CERT), [UDP Amplification Attack](#)
4. 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.

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.

APPENDIX A – How NetForecast Reports Meter Accuracy

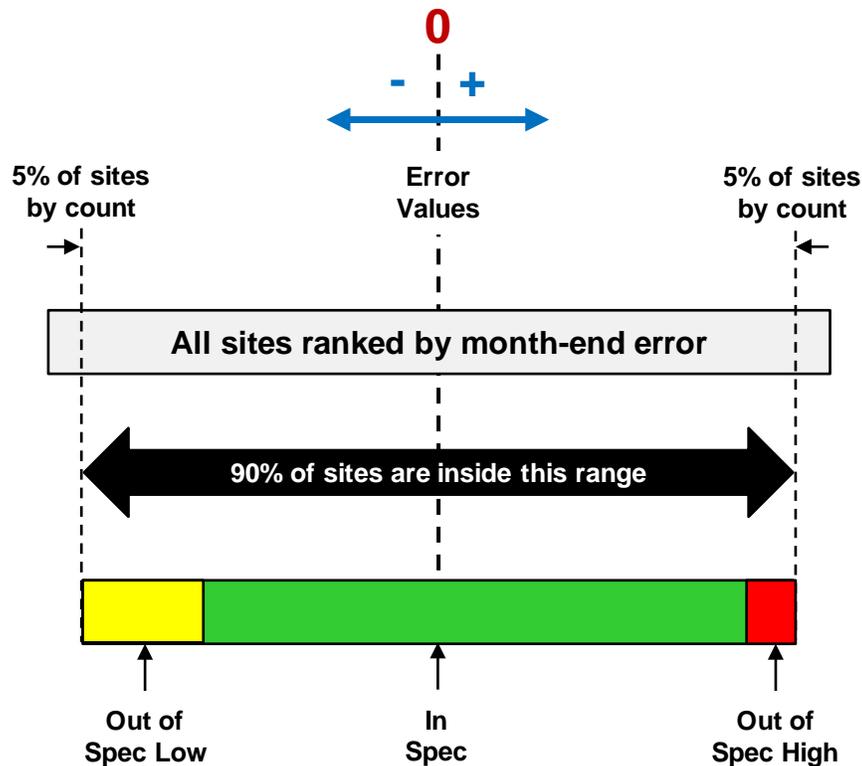
The following is a general high-level description of the NetForecast meter accuracy reporting methodology. Descriptions and examples in this appendix are based on a +/-1% error specification.

Defining the Monthly Dataset

Assume there are 50 month-end values per month (one per site). Averaging is not informative. It is important to examine the two edges of accuracy: How far was negative error from zero and how far was positive error from zero?

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

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



All subsequent error analysis and summaries are performed on the middle 90% of the dataset as shown above denoted by the black arrow. The 5% trimming is not applied in months with fewer than 20 sites in the results dataset. Under those conditions all results are used in the evaluation.

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*.

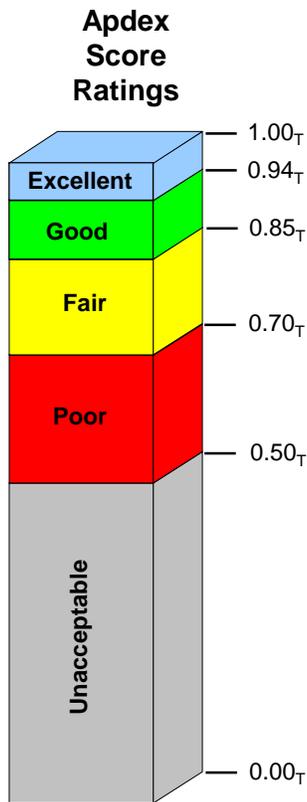
The two views are analyzed in three levels of detail as shown in the table below. Results with an “X” are documented in this report.

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

High-Level 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 vendors and enterprises 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% (under-counting). Under counting is outside the specification; however, some modest under-counting can be tolerated since the subscriber is not harmed by some traffic not being counted.
- Incorrect*: Number of samples above +1% (any over-counting). Any over-counting above the 1% specification is a serious problem that users will not tolerate.

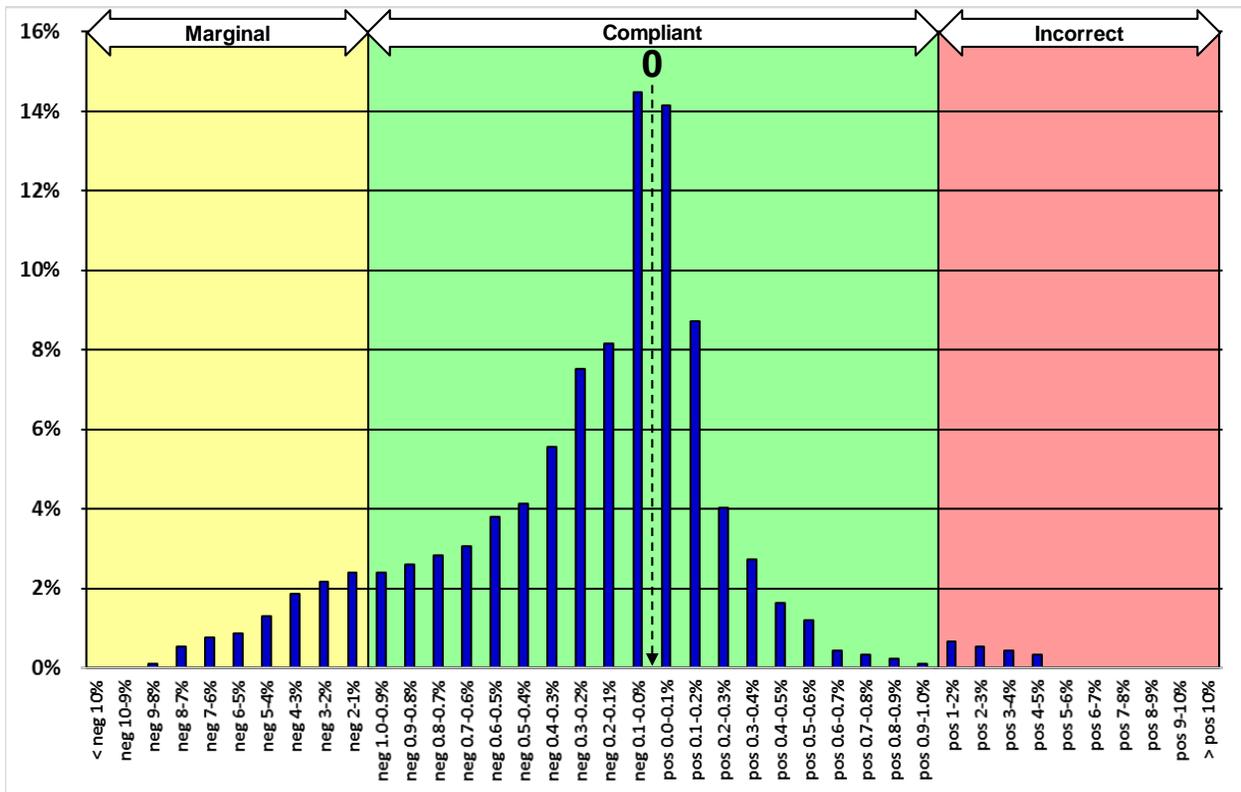
The Apdex methodology converts many measurements into one number on a uniform scale of 0 to 1 (0 = completely inaccurate; 1 = perfect accuracy). The resulting Apdex score is a numerical measure of accuracy performance. The Apdex formula is the number of compliant results, plus half of the number of marginal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

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

Incorrect samples have zero value

Target T or specification is always shown with the score

The chart below shows a typical meter error distribution or frequency plot. Cumulative daily error distributions and month-end error distributions show a similar pattern. In the example below, the samples highlighted in green are within a +/-1% specification, while samples in yellow are out-of-spec low (under-counting) and those in red are out-of-spec high (over-counting). 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 under-counting: 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: $\text{Apdex} = 0.88 + 0.10/2 = 0.93$. The distribution has an Apdex score of 0.93, which has a rating of Good.

The Apdex methodology is applied to both cumulative daily and month-end errors each month. The overall Apdex score and rating for the reporting period is the average Apdex score and its associated rating across the reporting period.

Summary results are shown in a table such as:

| | Cumulative Daily Error | Month-End Error |
|--------------|------------------------|-----------------|
| Apdex Score | 0.88 | 0.95 |
| Apdex Rating | Good | Excellent |

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

The frequency with which sites fell within or outside the +/-1% accuracy goal each month is determined on both a month-end basis. The following table shows the percentage of sites: within the specification, out-of-spec low, and out-of-spec high.

Out-of-spec high: Identifies the single worst out-of-spec high month.

Within spec: Shows the average site percentage within spec across the months covered by the reporting period. This is the overall probability that sites were within specification.

Out-of-spec low: Identifies the single worst out-of-spec low month.

| | Metric | Percent of Sites In/Out of Spec | Month |
|-------------------------|-----------------------|--|--------------|
| Out-of-Spec High | Worst month | 10% | January |
| Within Spec | Average of all months | 80% | NA |
| Out-of-Spec Low | Worst Month | 20% | July |

Note that the percent of out-of-spec sites may not add up to 100% because the results in each row are not from the same month.

Meter Error – How Far the Meter Is Outside the Specification

The final and most detailed view shows the average magnitude of the meter errors within each error group during the measurement period.

Out-of-spec high: Identifies the average error in out-of-spec low sites during the worst out-of-spec month. Worst month here is defined by the average error, which may or may not correspond to the worst out-of-spec month as defined above.

Within spec: Shows the average error for sites falling within the specification across the reporting period. This is typically a value close to zero since it is an average of all sites over all months that were within the spec.

Out-of-spec low: Identifies the average error among out-of-spec low sites during the worst out-of-spec month. Worst month here is defined by the average error, which may or may not correspond to the worst out-of-spec month as defined above.

| | Metric | Average Error Value | Month |
|-------------------------|-----------------------|----------------------------|--------------|
| Out-of-Spec High | Worst month | 4% | January |
| Within Spec | Average of all months | 0.1% | NA |
| Out-of-Spec Low | Worst Month | -8% | July |

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 over-counted. Negative error indicates the relative amount the ISP under-counted. |
| 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%). |