



Accuracy Assessment of Suddenlink Communications' Internet Usage Meter

By Peter Sevcik
April 2013


Suddenlink Communications has deployed a data usage meter system to measure and report how much Internet data traffic a subscriber consumes and generates. Suddenlink reports Internet subscriber usage information on its customer portal. Suddenlink engaged NetForecast to perform an independent audit of the accuracy of its data usage meter system.

NetForecast generated test traffic, performed independent traffic measurements, obtained daily usage meter records from Suddenlink for the test locations, and compared NetForecast's data with Suddenlink's records. NetForecast performed tests in Tyler, Texas and Malvern, Arkansas during a six week period in February and March, 2013. Suddenlink had no prior knowledge of the timing or volume of NetForecast's test traffic.

Suddenlink established an accuracy goal for its Internet usage meter to correctly measure traffic passing through a subscriber's cable modem within plus or minus (+/-) 2.0% over the course of each day. Our analysis validates that the meter is accurate within plus or minus (+/-) 0.8% over each day at each location—an outcome that falls well within Suddenlink's 2% stated goal.

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 period (see Fig 1), by month (see Figure 2), and by day (see 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.



NetForecast Report
NFR5115

©2013
NetForecast, Inc.

Figure 1 – Sample Subscriber Summary View of Usage

Internet usage reports are available via Suddenlink's customer portal at <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 at suddenlink.net, selecting **My Account**, then **Online Tools**, and then **View Internet Usage** to obtain the summary view shown in Figure 1.

The monthly and daily views shown in Figures 2 and 3 are available from a clickable list on the left of each view. The daily usage view has a pull-down menu to select previous billing months with up to 6 months of history.




Figure 2 – Sample Subscriber View of Usage Reporting by Month




Figure 3 – Sample Subscriber View of Usage Reporting by Day¹

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 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 two test sites were supported by a Cisco

¹ The usage views shown in Figures 2 and 3 were taken from different sites.

uBR7000 or uBR10000 series CMTS. A CMTS periodically reports the down and upstream counts in an Internet Protocol Detail Record (IPDR) as shown in Figure 4.

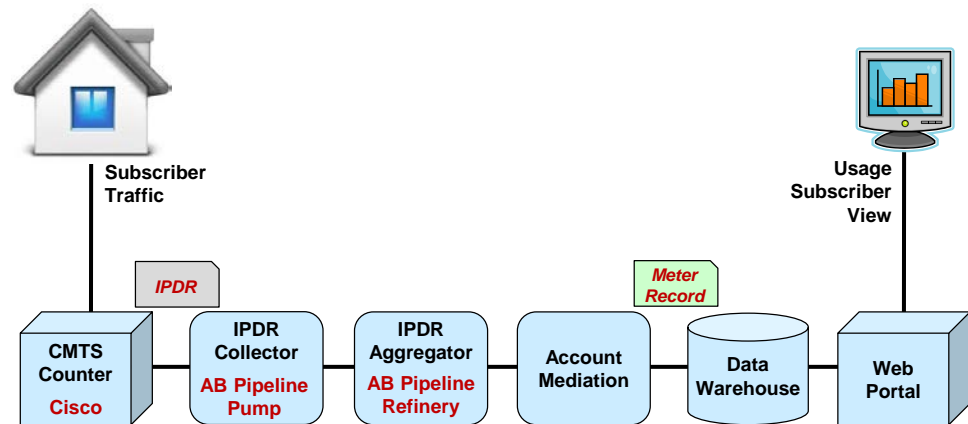


Figure 4 – Usage Meter Process Flow

From the CMTS, IPDRs are sent to an Applied Broadband (AB) 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

Suddenlink defined specifications for the accuracy and functionality of its Internet usage meter. NetForecast independently validated the meter relative to Suddenlink's stated goals for: traffic counted, data accumulation period, error bounds, timeliness, exception handling, granularity, mathematical consistency, accessibility, availability, and clarity as defined in the meter specification. A comprehensive description of usage meter specification factors is available in NetForecast's report, *ISP Data Usage Meter Specification Best Practices for MSOs*, reference [1].

As the usage meter system auditor, NetForecast has no stake in the design or implementation of the Suddenlink usage meter, nor do we endorse the usage meter specification. The objective of our testing is to assess whether or not the system complies with the meter specification.

NetForecast obtained subscriber accounts and performed usage meter accuracy validation testing for all test locations. Both downstream and upstream testing were performed under a variety of conditions.

NetForecast instrumented the accounts with a test laptop PC running NetForecast software and a Linux-enabled NETGEAR router with special firmware. In addition, we used FTP accounts on various NetForecast servers on the Internet. Figure 5 shows the NetForecast instrumentation. A more detailed description of the NetForecast methodology is available at reference [2].

The NetForecast test involved scripts performing an FTP file transfer from one of our test servers to the laptop. The tests consisted of repeatedly transferring files of varying file sizes in complex patterns. These tests were also performed as uploads from the laptop to the server. The script generated a log file documenting the transfer results and capturing detailed timing information for each transfer.

Extreme care was taken to ensure that only test traffic was sent or received through the cable modem. The laptop and router were the only devices connected to the cable modem. The laptop was cleansed of all applications that could generate traffic not needed for the tests—and applications were configured to neither request nor receive any software updates. Finally, remote management of the laptop was scheduled not to occur during testing.

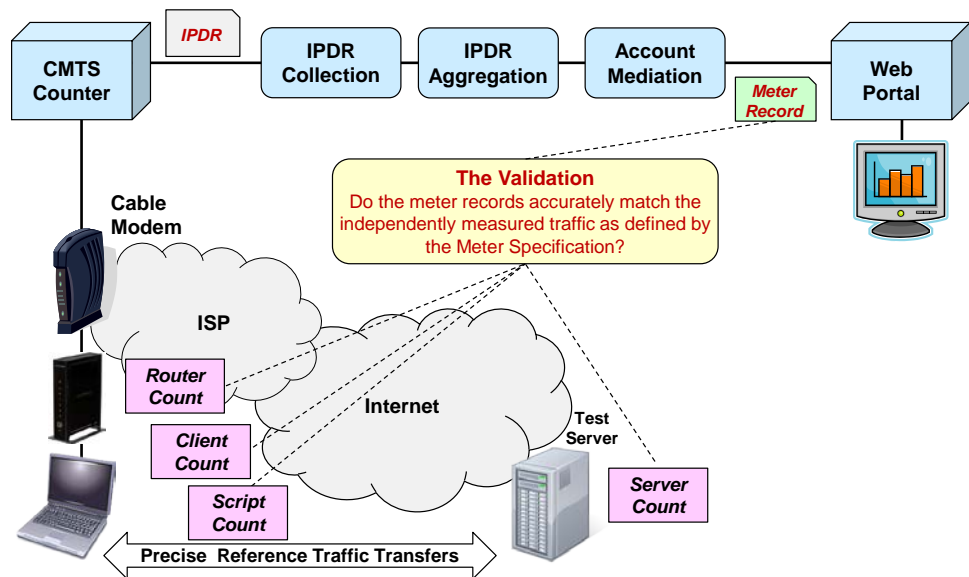


Figure 5 – The NetForecast Methodology

For each test, NetForecast produced four of our own carefully documented records of up and down traffic, illustrated by the pink boxes in the lower left of Figure 5. These record sources included:

- File transfer success logs on the test PC
- Detailed network traffic information as seen at the PC
- Upstream and downstream traffic processed by the router firmware
- Server records generated using detailed traffic instrumentation software.

Because of their locations in the network, each of the NetForecast measurement points sees different amounts of protocol overhead. Therefore, all four record sources were adjusted to account for upstream and downstream protocol overhead as counted at the CMTS.

NetForecast also gathered usage meter data from two places in the Suddenlink meter system during testing phase as shown in Figure 5. The data sources were:

- Daily traffic records from the data warehouse, which provided a preview of measurements presented on the Suddenlink meter portal
- The account portal for the test location, which reported the usage meter data to the subscriber.

Usage meter system accuracy validation entailed comparing data from the two Suddenlink reporting sources to the four sets of NetForecast-generated measurement data. Comparing the NetForecast and Suddenlink data enabled an end-to-end view of the Suddenlink usage meter's accuracy.

The NetForecast hour-by-hour data were seen by NetForecast only and were not shared with Suddenlink. Suddenlink did not know which tests NetForecast performed or when they were performed.

NetForecast Independent Suddenlink Usage Meter Validation Results

This report section summarizes NetForecast's validation test results for Suddenlink's subscriber Internet usage meter as seen from the subscriber perspective in Tyler, Texas and Malvern, Arkansas. NetForecast validated Suddenlink's specified goals for the usage meter against what we actually found.

1. What the meter counts

Suddenlink Goal: The Suddenlink Internet usage meter will count resource consumption on the subscriber's access line and within the Suddenlink network to/from the Internet. The Suddenlink usage meter will count subscriber traffic in bytes (eight bit increments also known as octets) and report the count in gigabytes on the subscriber portal.

All subscriber traffic to and from the Internet (known as payload), including Suddenlink management and voice over IP service traffic, will be counted, along with protocol overhead required to transfer the payload over the cable modem access line. *[Note that this overhead is contributed by a number of network protocols including: DHCP, DNS, Ethernet, IP, and TCP. These protocols are included within the error calculation discussed on page 6. The overhead varies by application, depending on the protocols the application employs to transmit user data over the Internet. A rule of thumb is that Internet protocols add from as few as 6% more bytes to the user data traffic for file transfers, to as many as 97% more for short packet-based traffic like VoIP.]*

Suddenlink will not include any of the DOCSIS protocol.

NetForecast Validation Result: NetForecast validated that the Suddenlink meter system counted all traffic except for DOCSIS overhead traffic.

2. Data accumulation period

Suddenlink Goal: The meter data accumulation period defines the minimum period over which the meter needs to gather a sufficient number of measurements to be accurate. Suddenlink's accuracy timespan goal is one day.

NetForecast Validation Result: NetForecast confirms a data accumulation period of one day.

3. Error bounds (Accuracy)

Suddenlink Goal: The error bounds define the acceptable error range goal for the meter. The error bounds goal for the Suddenlink meter system will be plus or minus (+/-) 2% over the one day data accumulation period (see above). A positive error means the meter over-reports subscriber traffic, and negative error means that the meter under-reports.

NetForecast Validation Result: NetForecast validated that the Suddenlink meter system was accurate to within +/- 0.8% over a data accumulation period of one day for the sites validated. This includes all the protocol overhead as defined in “What the Meter Counts” above.

NetForecast gathered 272 samples from the two cities tested. The overall error range across both tested cities was within +/- 0.8% over a day. The reporting system maintained that accuracy throughout all the elements of the Suddenlink meter system—up to and including the final view as seen on the customer portal. This statement applies when comparing the meter with the actual traffic sent to and from the cable modem, which includes protocol overhead as explained above.

Figure 6 shows the error distribution across the two test cities. A negative error indicates that the Suddenlink meter value is low relative to the NetForecast reference value (under reporting). A positive error indicates over reporting. All of the tests fell within the plus or minus (+/-) 0.8 percent range. The overall mean for all tests was slightly negative, indicating a slight bias toward under-reporting.




Figure 6 – Suddenlink Meter Error Distribution

4. Timeliness

Suddenlink Goal: The Suddenlink usage meter system will update the customer portal view every 24 hours. The goal is to update the customer portal each day at or before 8:00 A.M. Central Time, to reflect usage through 11:59 P.M. local time of the previous day.

NetForecast Validation Result: NetForecast confirms that the customer portal was updated as specified during the test period.

5. Exception handling

Suddenlink Goal: If a new update shows no incremental usage in the 24 hours since the last update, Suddenlink explains to subscribers that one of two conditions apply: “either (a) there was no incremental usage in that time period; or (b) the transmission of daily usage (from our counter to this page) was delivered too late to be deemed useful to the customer and was thus not counted and not applied to the customer’s usage total.”

NetForecast Validation Result: NetForecast confirms that the exception handling explanation is in place on the customer portal.

6. Granularity (How the data is shown)

Suddenlink Goal: The information received by the Suddenlink usage meter system will be in binary bytes, reported to the level of one-tenth gigabyte. Traffic measurements will be displayed to one decimal point, and the second decimal point (one-thousandth GB) will be rounded down (i.e., truncated), which means the numbers to the right of the decimal point are not shown (as seen in Figures 1, 2 and 3).

NetForecast Validation Result: NetForecast confirms that the data is displayed in binary bytes, reported at the gigabyte level, displayed to one decimal place (tenth) rounded down.

7. Mathematical consistency

Suddenlink Goal: Various views of a subscriber’s meter will provide a consistent traffic consumption value(s). The total usage count for each month is based on truncated values for each day. This means that the sum of values shown for each day in a month will match the value shown for the corresponding month value.

NetForecast Validation Result:
NetForecast confirms that daily and monthly usage counts match.

8. Accessibility

Suddenlink Goal: The Suddenlink usage meter reporting screens will be reachable within three clicks of the Suddenlink home page.

NetForecast Validation Result: NetForecast confirms that the usage meter is accessed from the Suddenlink user account home page as specified.

9. Availability

Suddenlink Goal: The Suddenlink online usage meter portal view will be available 99% of the time.

NetForecast Validation Result: NetForecast confirms that the portal view was available when accessed.

10. Clarity

Suddenlink Goal: The Suddenlink usage meter will be explained in a FAQ accessible at <http://www.suddenlink.com/allowanceplan/>.

NetForecast Validation Result: NetForecast confirms that the FAQ page is clear and informative.

Useful Information if You Want to Do Your Own Testing

If you wish to perform your own Internet Usage meter validation testing, it is important to be aware of factors that may cause your measurements to vary from what the meter reads. One such factor is where you measure. If you measure from the network, you will see protocol overhead that you will not see from a computer. Another factor affecting measurements is packet loss, and another is the presence of “unexpected” traffic.

Avoiding Binary versus Decimal Confusion

The Suddenlink meter reports in gigabyte increments. 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. This is important to understand because counting traffic by billions of bytes will result in an apparent (but not actual) 6.9% error in the meter.

Binary		
KB	Kilobyte	1,024
MB	Megabyte	1,048,576
GB	Gigabyte	1,073,741,824
TB	Terabyte	1,099,511,627,776

Decimal	
Thousand (Kilo)	1,000
Million	1,000,000
Billion	1,000,000,000
Trillion	1,000,000,000,000

Decimal and binary values are often mixed. For example, 10MB means 10 times 1,048,576 bytes. A display showing 1,000MB is actually stating 1,048,676,000 bytes, which is still not a GB as shown in the table above. It requires 1,024MB to reach the value of a GB.

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). This may falsely lead you to believe that the meter is over counting.

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 the user to erroneously conclude that the Suddenlink meter is over-reporting.

If you measure traffic in the network, you will see the payload traffic plus overhead from protocols like TCP/IP and Ethernet, which generally add up to 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.

Packet loss adds to the meter

Any reliable end-to-end protocol (like TCP) has a mechanism to retransmit packets lost in transit. Packet loss will add to the traffic seen by the meter in some situations.

“Unexpected” Traffic

If you look closely at your Internet usage, chances are you will see unexpected bytes register on your meter. Most Internet users don't know how much traffic their household produces. Here is a sampling of traffic sources that may surprise you.

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 just to PCs, but 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, and for convenience, these automated connections are invisible to the user so you may be unaware of the traffic generated.

In rare cases a PC could be hijacked and generate traffic that has nothing to do with any user in your home. Also, a neighbor or "wardriver" may use your connection without your permission or knowledge.

The most likely source of unexpected traffic, however, is PC software. Most popular software has automated update features which 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 aren't 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 sent overnight to a smart phone or MP3 player to be viewed or listened to during the morning commute.

Assume each night's upload 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 from Amazon or Netflix, etc.—and you may get more bits than

you pay for 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].

Tracking down rogue traffic

If you cannot account for a high traffic volume on the meter and suspect some rogue consumption, we recommend performing a controlled test. Plan for a solid period of time when the home can become "digitally silent" (e.g., overnight or on a weekend when traveling). At the start of the silent period, turn off all devices that can access the Internet. Make sure, however, to keep the router and cable modem operating.

At the end of the digital silence, turn on one PC and log back into the Suddenlink usage meter portal, or check from an Internet cafe or other means while you are away. Look at the daily traffic views for the time period that you observed digital silence. If true digital silence was achieved, the meter should not have incremented by more than 0.1GB in any day. If there is significant traffic shown, then there is certainly some other traffic consumer connected through the router.

Conclusions

The Suddenlink subscriber Internet usage meter is accurate to within plus or minus (+/-) 0.8 percent over the day and in aggregate over the month. The reporting system maintains that accuracy throughout all the elements of the meter system—up to and including the final view as seen on the customer portal. This statement applies when comparing the meter with the actual traffic sent to and from the cable modem, which includes some protocol overhead as explained above. Based on our test results, subscribers should be able to rely on the meter's accuracy.

NetForecast found the Suddenlink usage meter online reports easy to understand and use. NetForecast found the online system was always available when accessed during the testing period. We also found the descriptive information about the usage meter to the site to be adequate and clear.

References

1. Sevcik, *ISP Data Usage Meter Specification Best Practices for MSOs*, NetForecast Report 5110 (Updated), February 2013.
2. Sevcik, *Validating the Accuracy of ISP Subscriber Traffic Usage Meters*, SCTE Cable-Tec Expo, Atlanta, GA, November 15, 2011.
3. 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 Author

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 two patents on application response-time prediction and congestion management. He works extensively with the SamKnows system in support of the FCC Measuring Broadband America project, analyzing operational integrity and performing deep data analysis. He can be reached at peter@netforecast.com.

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.