> Itron white paper

Time Accuracy in the 100G Datalogging Endpoint

Chris Osterloh
Sr Principal Systems Engineer

Itron

Executive Summary

Itron's 100G Datalogging endpoints are designed to deliver time synchronized interval data. Support for this capability is provided in endpoint model part numbers ERG-5002-XXX and ERG-5003-XXX. Time synchronized interval data is accomplished by the use of a real time clock (RTC) in each endpoint. This paper describes the operation and accuracy of the clock in both synchronized and non-synchronized environments.

100G Datalogging Endpoint Time Basics

Every 100G Datalogging endpoint contains a RTC tracking seconds since midnight. This clock is used to control interval data collection and other operations such as Gas Day Take processing.

As the RTC in the endpoint passes through the top of the hour, a record of consumption for the last 60 minutes is recorded in a log with 960 total entries (24 hours a day for 40 days). A log representation is shown in the table below. The blue section contains the current meter consumption at any instant in time. It also contains the partial amount of gas used in the current unfinished hour. These values are updated each time the endpoint records any new consumption. The yellow sections contain the historical hourly gas amounts consumed over the last 40 days (960 hourly intervals). These amounts do not change once calculated.

Note: The first entry is the most recent record in the log and the 960th is the one made 40 days ago. As the end of the current hour interval is completed, the partial hour becomes the new first entry. The old first entry becomes the second entry, the 959th entry becomes the 960th and the old 960th is removed.

Consumption	123456
Partial hour (interval #0)	3
1	12
2	3
3	2
4	5
:	:
:	:
960	2

From this data set, we can determine four primary things:

- Current meter consumption (the index read)
- Consumption at the end of the last hour (current consumption – partial hour)
- Delta Consumption over any 60 minute period over the last 40 days
- Consumption reading at any time in the last 40 days (current consumption – partial hour – all intervening intervals)

There are two primary operating environments for the 100G with respect to time. A synchronized environment that allows for a worse-case accuracy under a network application of ±1 minute and a non-synchronized environment that allows for a worse-case accuracy of ± 30 minutes. These two environments are described in more detail in the remainder of this document.

100G Datalogging Endpoint in a Time-Synchronized Environment

The most accurate mode of operation for the 100G Datalogging endpoint is in a time-synchronized environment. This means the reading systems monitor and regularly synchronize the RTC in the endpoint. Fixed Network, Mobile and Handheld solutions all support this type of operation.

Fixed Network Operation (CCU 100)

The time is initially set in the 100G Datalogging endpoint when it is first programmed, which is typically during installation. After that point, each data message sent up from the 100G contains the time from the clock in the endpoint. The CCU 100 Collector compares the endpoint time with the master clock in the collector. If the time is off more than a specified amount (± 45 seconds is the default), the CCU 100 sends a time-synchronization command to the endpoint. Upon receipt of the time-synchronization command, the endpoint resets its internal clock and shortens or lengthens the current interval by the appropriate amount. The interval where the time correction occurs is marked as corrected.

In Itron's Gas AMI solution, collectors are calibrated to coordinated universal time (UTC) using an integral GPS receiver as the primary time reference. GPS-derived UTC allows the collectors to operate as if they are continuously connected directly to a stratum 1 timeserver. GPS derived UTC is directly traceable to NIST and varies less than 100 nanoseconds from the NIST master clocks.

In the unlikely event of an extended loss of GPS signal, a Fixed Network Collector generates an alarm and defaults to synchronization through Network Time Protocol (NTP). Itron recommends using a configurable, stratum 1 NTP server located at the head-end. Due to differing backhaul options and varying network transit conditions, this link normally maintains a time differential with NIST of < one second. If a collector is not synchronized to a master UTC reference through either GPS or NTP, the collector is not allowed to set or correct endpoint times.

All interval data collected by the endpoint is referenced to its internal clock. As such, and with the processes presented above, this allows the Itron Gas AMI solution to provide time-synchronized interval data from the 100G Datalogging endpoint that is within one minute of a NIST traceable source when operating in a fixed network.

Mobile and Handheld System Operation (MC3, MCLite, FC200SR or FC300 with SRead)

As with the Fixed Network solution, the time is initially set in the 100G Datalogging endpoint when it is first programmed, which is typically during installation. After that point, the synchronization of the endpoint's RTC is automatically performed every time data is requested from the endpoint (default reading system configuration). For endpoints that are time synchronized on a regular basis, the maximum RTC drift is four minutes per month. For example, if the utility requests 40 days of daily data from each endpoint every month, the endpoints will remain time-synchronized to within 4 minutes of the collector time. Itron systems allow for automatic time synchronization between the handheld and the server. For mobile collection systems, Itron recommends the utility establish a process for accurately setting the mobile collector laptop's time.

100G Datalogging Endpoint in an Non-Synchronized Environment

When the 100G is operating in a non-synchronized environment, the reading system is not setting the RTC (RTC) on the endpoint on a regular basis. In this mode, the clock will drift over time but the architecture of the data collection system will limit the error to ± 30 minutes maximum.

In this application, the meter reading radio still contains a RTC, which is set to accurately reflect the current local time (for example, 12:05 p.m. March 15, 2010). In a traditional meter reading system, it is the reading device which would normally collect the consumption data from the endpoint and add a time stamp to this reading. The consumption value, when read by the Itron collector, becomes the consumption at a certain time.

The ability of the reading device to add the actual time stamp to a meter read is very important to the concept of time in the 100G endpoint, as it is the collector time tying endpoint consumption to a specific time in a non-synchronized environment.

As discussed in the previous section, the endpoint contains a log of the current total consumption and the last 40 days of hourly intervals. For now, let's assume the time in the collector and the endpoint correctly match each other.

In this mode, the endpoint only cares about the minutes in the hour (05), and the collector contains the actual "real time" (12:05 p.m., March 15, 2010). From the collector's view, the data in the meter log can now be aligned with the actual real time. The collector time is shown in the gray column of the table as a reference.

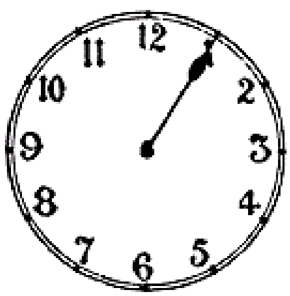

Endpoint

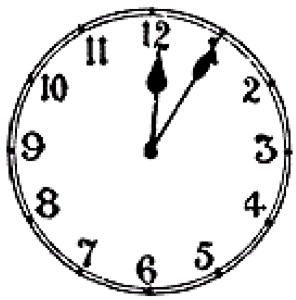

Collector

Current consumption	123456	12:05 p.m. March 15, 2010
Partial hour (interval #0)	3	12:00 p.m. - 12:05 p.m. March, 15, 2010
Interval #	Value	
1	12	11:00 a.m. - 12:00 p.m. March 15, 2010
2	3	10:00 a.m. - 11:00 a.m. March 15, 2010
3	2	9:00 a.m. - 10:00 a.m. March 15, 2010
4	5	8:00 a.m. - 9:00 a.m. March 15, 2010
:	:	
:	:	
960	2	12: p.m. - 1:00 p.m. February 4, 2010

The collector can therefore request consumption from the endpoint at any time from February 4 to March 15, 2010. It only needs to request that the endpoint go back the required number of hours.

For example, if the collector requires a consumption at 9:00 am March 15, 2010, it will request the endpoint go back 3 hours. The endpoint will then take the current consumption (123456) and subtract the amount consumed between now and 9 a.m. (3 + 12 + 3 + 2 = 20) and return the consumption reading at the requested time (123436). This is the basis of the time computation in non-synchronized environments for the 100G endpoint. In a non-synchronized system (one where time has never been set), the above approach will always yield a measurement accurate to within one hour. However, further improvement is achieved using the rule of ½ hour described in the next section.

The Rule of ½ Hour and What Interval to Return

As previously described, when operating in a non-synchronized environment, the endpoint device uses the minutes in one hour only. The following example illustrates how a ± 30 minute maximum error in requested interval data is achieved even when the system is not time-synchronized.

For example, suppose the endpoint time is at 58 minutes and the collector is at 12:59 p.m., the clocks have only a one-minute difference. The collector wants a reading from 10:00 a.m., so it asks the endpoint to go back to interval 2. The partial interval is interval 0, the 11:00 am to 12:00 noon block is interval 1 and the 10:00 a.m. to 11:00 a.m. block is interval 2. The endpoint then calculates (123456 – 3 – 12 – 3) and returns (123438) as the correct consumption at 10:00 a.m.

If the collector were to read the same endpoint one minute later, the endpoint is at 59 minutes, and collector is at 1:00 pm. The collector now tells the endpoint to go back to interval 3. Since the RTC in the endpoint is still in the past hour, the partial interval is interval 0, the 11:00 a.m. to 12:00 p.m. block is interval 1, the 10:00 a.m. to 11:00 a.m. block is interval 2 and the 9:00 a.m. to 10:00 a.m. block is interval 3. So the endpoint then calculates (123456 – 3 – 12 – 3 – 2) and returns (123436) as the requested consumption, which is actually the 9:00 a.m. reading, not the desired 10:00 a.m. reading.

Although there is only a difference of one minute in the clocks, there is one-hour uncertainty in the data. In this example, the collector would have to look at the returned data and the endpoint time, and then re-request the data, using one less hour to get the correct reading. This adds extra time to a read, and makes the collector much more complex, as it now needs to evaluate the returned data to see if it is correct.

The solution to the problem is simple and lies in the rule of ½ hour.

In the above example, suppose the collector sends its time along with the request for data specifying a number of intervals back. The endpoint evaluates the rule of ½ hour, and determines there is more than 30 minutes difference between the collector time and the endpoint time. This provides a key additional fact – namely the hour requested is not what the collector wants. This additional fact allows the endpoint to change the data requested by ±1 hour and guarantee the data returned is the closest interval to the one the collector wanted. This adjustment negates the need to re-request data in a time-synchronized system, and guarantees the data returned in a non-time-synchronized system is within ½-hour accuracy. Since most users want to resynchronize their time on a monthly basis, it does not add overhead to the protocol, as time setting is accomplished at the same time; if the endpoint knows the collector time, it can pick the best interval to return.

The rule of ½ hour is implemented as follows:

- When the endpoint device receives a request for additional data, a check is performed to see if collector time is included.
- If collector time is not included, the response for the specific hour requested is returned.
- If collector time is included, the endpoint device extracts the collector time within the hour (0 – 59 minutes).
- If the absolute value of the difference between collector time within the hour and endpoint time within the hour is less than 30 minutes, the endpoint returns the specified hour.

- If the collector time – endpoint time in the hour is < -30 minutes, then the endpoint adjusts the time requested by one less hour. If one less hour results in an hour value of 0, then an invalid interval result is returned.
- If the collector time – endpoint time in the hour is > 30 minutes, the endpoint adjusts the time requested by one more hour. If one more hour results in an hour value of 961, then an invalid interval result is returned (the endpoint device stores 960 intervals).

An Itron mobile reading device (i.e. MC3, MCLite, FC200SR or FC300 with SRead) in conjunction with a 100G Datalogging endpoint module, when operated in a non-synchronized time environment, has a worst-case stability of ± 30 minutes due to the ½-hour rule.

Conclusion

100G Datalogging endpoints are designed to provide accurate time-stamped hourly interval data. When 100G DL endpoints are operating in a time-synchronized Fixed Network environment, the maximum time deviation is ± one minute, while mobile and handheld environments allow for a maximum ± four minutes per month of drift. This provides unparalleled data accuracy for gas utilities utilizing Itron's 100G Datalogging endpoints in either a fixed or a mobile environment.

About Itron

At Itron, we're dedicated to delivering end-to-end smart grid and smart distribution solutions to electric, gas and water utilities around the globe. Our company is the world's leading provider of smart metering, data collection and utility software systems, with nearly 8,000 utilities worldwide relying on our technology to optimize the delivery and use of energy and water. Our offerings include electricity, gas, water and heat meters; network communication technology; collection systems and related software applications; and professional services. To realize your smarter energy and water future, start here: www.itron.com.

Itron Inc.
Corporate Headquarters
2111 North Molter Road
Liberty Lake, Washington 99019
U.S.A.
Tel.: 1.800.635.5461
Fax: 1.509.891.3355

Publication 101106WP-01
12/10