

Leap Second Readiness Tips: A Guide for Managers of Critical IT Infrastructure



1: Decide If You Should Care

Besides impressing your friends, what should you do with the knowledge of an upcoming leap second at UTC midnight on December 31, 2016? If your network operations do not currently rely on synchronizing events, then you probably don't need to care. Time-stamps across different machines are probably off by more than one second already. You can skip right to tip #7, then consider how synchronizing your network may help with troubleshooting, security, and overall efficiency of your network applications. If, however, your systems are synchronized you need to be aware how your infrastructure, systems, and applications will handle the upcoming leap second.

2: Take an Inventory of Your Time Servers

This may be easier said than done, since time synchronization deployments, particularly those based on NTP, sometimes evolve organically rather than being highly architected and documented. But if time synchronization is important to you, then you will want to know all "time masters" used by your network elements. Do you know which time servers are synchronizing your systems? Are they all within your network perimeter? What are they syncing to?

Ask us about tools to automatically survey your network to track down the source of time server packets within your network infrastructure.

3: Evaluate if Your Time Servers are Leap Second Ready

If you identify a single time server, or a combination of time servers operating from within your network, check with the supplier(s). If they are network appliances, check for leap second readiness information from the manufacturer. If it is a software-based time server, check your current version, and any current release notes and documentation for leap second handling. Look for updates relating to leap second handling.

4: Check the References of Your Time Servers

In addition to what you may learn when working through tip #3, there are a few configuration items to pay close attention to. The most important is the identification of the time reference(s) for your servers. Many network synchronization deployments use the GPS satellite signals as their primary reference. This is a good time to check to make sure those time servers are currently operating with good GPS reception. You can check that the time server has noted the December 31, 2016 leap second indicator within the GPS broadcast, and that the proper GPS-to-UTC offsets are noted. If they are not, then you can remediate any problem with the GPS reception and leap second information handling.

If a reference other than GPS is used, then you need to analyze its incoming data stream. There is a long legacy of precision time

sources which stream time codes to slave devices. However, some time codes transmit information about leap seconds and some do not. Determine the ability of the upstream device to know of an upcoming leap second, how it transmits that information, and how the slave devices makes use of the information.

5: Check a Few Other Configuration Parameters

If you have come this far, then you should ask if there are any other time server configurations that deviate from standard leap second handling as specified by the sync protocols such as NTP and PTP. For example, some devices may allow the use of a time scale other than UTC. If the time server is configured for another time scale like GPS or TAI, then make sure you know why.

6: Now for the Tough Part – Understanding Your Operating System and Time Used in the Application

Even if all questions to this point have pointed to good leap second handling, the “moment of truth” is not revealed until the local synchronization client decides what to do on UTC midnight of December 31, 2016.

There are only 3 scenarios for what happens to time in the OS/ application:

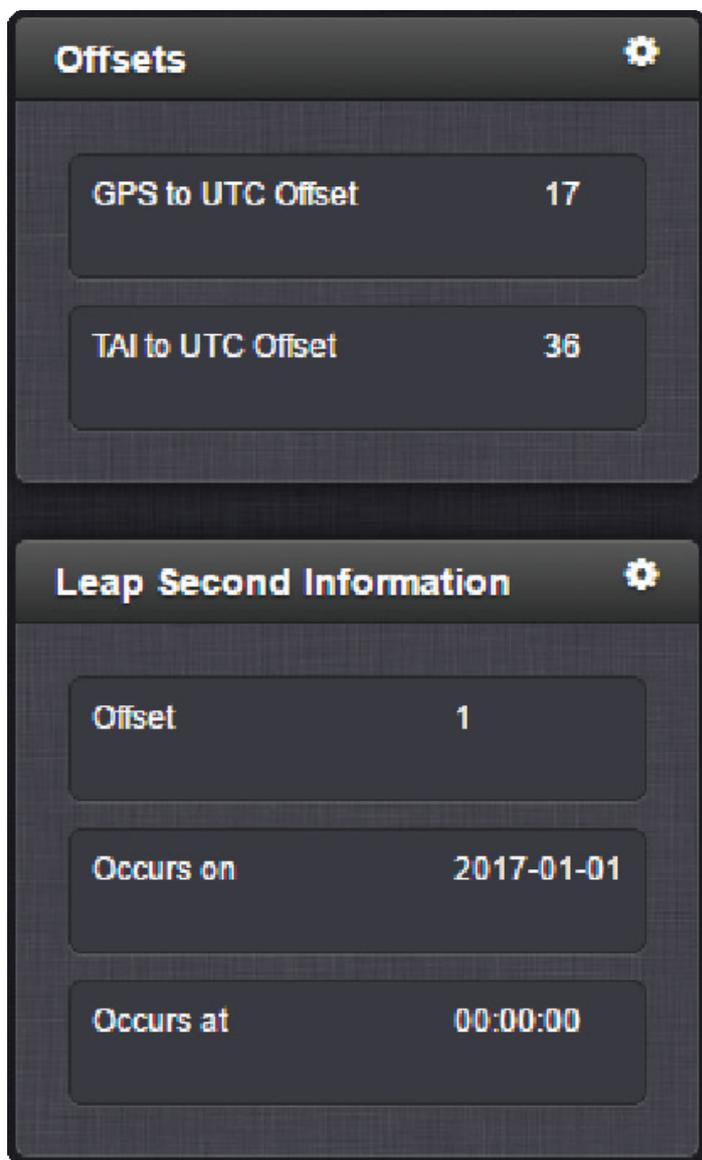
1. Time can freeze for the duration of the leap second.
2. Time can jump back 1 second.
3. The time increment for each “tick” can be greatly reduced at the leap second event to keep time monotonically increasing but converging to realign with UTC after the leap second event.

It is not actually that important to know what happens in all of your client implementations. What matters is the ultimate effect on your application. We know during previous leap second events computer systems were shut down and web servers crashed. A test network can be configured to understand the effect of a leap second and may be the only way to know for sure what will happen.

Ask us about leap second testing with any combination of a configurable time server, or a GPS simulator.

7: What are Your New Year's Plans?

Unlike the last leap second, this one will occur outside normal business hours. December 31st falls on a Saturday and most of the world celebrates the new year. However, network operations are expected to run 24/7. Check the time of the leap second event in your locality and plan to be ready to implement remedial actions in case the leap second creates a problem.



A Spectracom time server showing the event time scale offsets and pending leap second information.