Securing Reliable, Secure Time Across Your Data Centre Estate

GPS, and now other GNSS sources such as GLONASS, Galileo and BeiDou, has been a major contributor to frequency and time sync in Tier 1 services in telecoms and enterprise for over 20 years. From the release of the first commercial NTP Server – the TymServe 2000 – to the latest generation is capable of serving over 300,000 NTP requests per second in a data centre.

©Chronos Technology

Most of these systems – whether a frequency reference for digital telecom networks or a PTP time and phase service for next generation telecom and enterprise networks – have a one-to-one relationship between the GNSS source and the clock.

Whether a hosted shared data centre or one hosted and dedicated to a single enterprise, these sites usually have roof space for GNSS systems along with rack space for dedicated systems. This leads to many rooftops hosting tens of GNSS antennas with often less than optimal spacing between them.

Distributing GNSS and Time in a Data Centre

One method employed by some data centres is to distribute GNSS signals across the data centre. However, this doesn’t allow for the installation of a customer’s own GNSS system.

These systems will tend to use two or more antennas for resilience. They usually distribute through a combination of fibre and coaxial cables to minimise losses.

This GNSS service will deliver a coaxial cable with appropriate connection in the customer’s clock(s). There will be a calibrated time delay between the roof and the rack to ensure the best possible time sync.

An extension of this service is the delivery of NTP, or more usually PTP. These are packets direct to the customer’s clock device or servers themselves. In many instances, such a service may be used as a secondary input to a clock system using GNSS as the primary source. However, in a Financial Trading setting, this time sync service can deliver straight into the Trading Engine itself.

In this case, especially in a shared data centre, there needs to be clear demarcation between the networks of the data centre and customer and is often a step that many security departments in high value applications will simply not take. There is always a risk that a compromise of the data centre’s network. For example, a UDP Port 123 attack in an NTP system, could crossover into customers' networks too, with no real security measures available to the customer to counter this.

These services though only really apply to single data centres, however large. So, what time sync security measures can be done across a network, and indeed across a country or region?

Data Centre GNSS Security

GNSS signals can be easily jammed or even spoofed to deliver the wrong time and/or position to the clock system. Spoofing is technically difficult and is usually targeted at Critical National Infrastructure. There are built in safeguards in modern GNSS receivers and some external systems, such as Microchip's BlueSky GNSS Firewall.

All these GNSS signals are very low power and are found in the same area of the radio spectrum. All it takes is a single jamming signal from a device to take out all these GNSS services. Modern clock systems have internal oscillators to holdover the clock performance should GNSS signals be lost. At least for the period it takes to mobilise and fix the issue. GNSS loss is usually due to a faulty component in the GNSS system or someone cutting the wrong cable.

ePRTC and vPRTC

The ITU has specified the Enhanced Primary Reference Clock (ePRTC) to add residence to GNSS failure for mobile networks needing microsecond time sync for effective operation.

The ePRTC uses Caesium Atomic System Clocks in conjunction with your PTP Clock system to deliver a time sync 'flywheel' that can give you virtual autonomy from the sky. Although rooted in telecoms, this concept is equally valid in enterprise and utility applications. Examples include use in Financial Trading and Digital Substations. After an initial training period of three weeks, an ePRTC will be able to hold time sync to within 100ns of UTC for at least two weeks should GNSS signals be lost.

This timing can be distributed across your data centre estate using High Performance Boundary Clocks at most locations. The safest way to prevent GNSS jamming is avoid needing GNSS at your location. Using a wavelength of a direct optical connection between your data centres can deliver sub 5ns of additional time sync error across distances up to around 100km. Even Carrier Ethernet connection can bring time transfer to the tens of nanoseconds.

Using these techniques, you can build what has been termed a Virtual Primary reference Time Clock (vPRTC). Using the optical techniques each Point of Presence can have a Tier 1 timing system too within 100ns of UTC.

Summary

Time has been a Cinderella service in most data centre settings. However, the requirement for highly accurate time sync across has never been more important.

A well-designed time dissemination network across your estate will deliver accurate and reliable time in a secure way. It will also minimise the security threat to GNSS signals at the same time.

This is of enormous benefit whether the data centres are purely for your own services, or if you are a host striving to maximise the security and reliability of series to your customers, timing or otherwise.