Network Time Server Guidance

Answers and recommendations for CISA Time Guidance for Network Operators, Chief Information Officers, and Chief Information Security Officers

The U.S. Cybersecurity and Infrastructure Security Agency (CISA) recently published a Time Guidance document filled with recommendations and probing questions regarding proper network timekeeping to bring attention to critical aspects of synchronizing time across a network[1].

This Network Time Server Guidance document goes the next step to show how a Microchip dedicated Stratum 1 network time server, inside the firewall and properly configured, answers the CISA questions conforming to CISA guidance. This document cites key sections in the CISA document (annotated with bold italics) where a Microchip SyncServer S600 Stratum 1 network time server brings a network into compliance and follows the organizational flow and themes of the CISA document.

Background: Why Time Synchronize

The most basic question to be answered is why synchronize the time on a network in the first place. The answer is that networked devices need synchronized time for log file accuracy, network security systems, email servers, phone systems, general network operations and directory services, legal and regulatory requirements, and the list goes on. In short, accurate, secure, and reliable time-of-day is mandatory for improved network and business operations.

Network Time Protocol (NTP) is by far the most popular network protocol in use to synchronize time on network connected devices. NTP uses a conceptually simple hierarchy to transfer time from one device to another. Time typically originates at a reference source such as the atomic clocks at the US Naval Observatory (USNO) or NIST. In the case of USNO, the time is transferred to the atomic clocks on board the GPS satellites and broadcast back to the planet for use, free of charge. These high accuracy, high precision, systems are called Stratum 0 time references. Other Stratum 0 sources of time could be a terrestrial radio broadcast (e.g. WWVB, eLORAN), or a time server physically connected to the atomic clock.

Time servers that get their time directly from the Stratum 0 reference are referred to as Stratum 1 clocks. The typical Stratum 1 clock is a network time server with a GPS receiver installed as it is getting its time directly from the GPS Stratum 0 reference. The stratum of any time server is determined merely by the stratum of the time server it is getting its time from, plus 1. Think of it as a clock hop count.

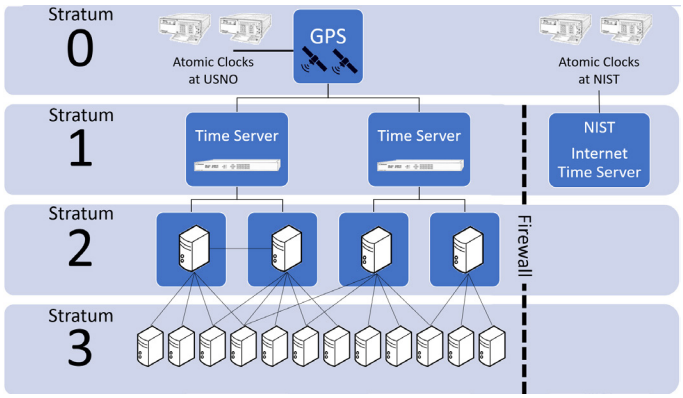

Figure 1. Time of day distribution through the NTP Stratum hierarchy

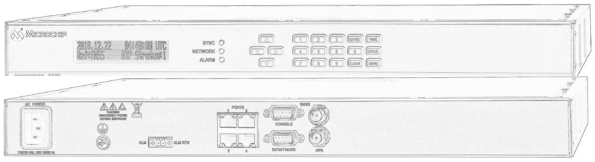

Figure 2. Typical Stratum 1 network time Server (front/back)

Stratum 1 time servers play a unique and essential role in the network time hierarchy as they are the first reference source of time connected to the network. These devices bridge the gap between every clock in the NTP network hierarchy and the Stratum 0 source of time. To that end, Stratum 1 time servers define the accuracy, reliability and security of time for all other strata of clocks beneath them and the applications that rely on those clocks.

1. Knowing Your Timing System

Traceability of time on your network is essential. While you may know where the time is coming from through the hierarchy inside the firewall, the chances are that unless you have a Stratum 1 time server inside the firewall you are likely going through the firewall to a time server on the open internet, perhaps even from a pool of time servers.

[1]https://www.cisa.gov/sites/default/files/publications/time_guidance_network_operators_cios_cisos_508.pdf

While this "*free time from the internet*" is convenient, it is a significant vulnerability in terms of accuracy, reliability, and security of the time your systems are dependent on. Afterall, it is not your clock, you have no idea if its accurate or where it gets its time, it will not send an SNMP trap if the time is wrong, it is subject to packet manipulation and denial of service attacks on the open internet, you don't know if it has been patched to keep it from being hacked, etc. Other than an IP address that responds to NTP time requests, and is easily known to every bad actor, you have very little idea about an internet time server. A Stratum 1 time server inside the firewall mitigates the magnitude of all these risks.

Accuracy to UTC is important for networked systems. In truth, at the core of every networked device, be it a Linux server, Windows machine, etc., is time being kept to UTC. An organizations time accuracy requirements will vary and it may be very detrimental if the time on the network is not accurate enough. For most enterprises, the accuracy of the log file time stamps are probably the singularly most important item as they play an essential role in network forensics when there's a time critical problem to be resolved. A Microchip SyncServer S600 Stratum 1 network time sever for example is accurate to UTC to around 15 nanoseconds, with NTP time stamps as good as 20 nanoseconds to UTC, and while that may seem like overkill, its not as you will learn further below when it comes to holdover requirements and futureproofing timing on the network.

Time distribution methods to achieve desired UTC accuracy on key systems has evolved over time. Decades ago, Stratum 1 servers had limited time stamp throughput and a traditional stratum hierarchy solved the problem at the expense of accuracy. Today, a SyncServer S600 can effortlessly timestamp 360,000 NTP client requests for time per second with timestamp precision of 1ns and 20ns accuracy to UTC. With the speed of LANs these days and the performance of Stratum 1 time servers like the SyncServer, flattening the NTP hierarchy and going direct to the SyncServer is a very viable option.

IT inventory of time servers and software updates are very important. As a leading manufacturer of time servers, we routinely hear stories of old Stratum 1 time servers that had been forgotten, not maintained, yet still provided time to the network. Admins did not know where they were located, only the IP address(s). The story always ends the same, the time server failed for some reason after many years of service, the network faulted in some way and panic set in. This has been the case for time received from the internet as well because local machines are often multipurposed to operate as time servers synchronizing with internet time servers. These machines are often not maintained over the years as IT personnel come and go and the ultimate decommissioning of the machine can result in the inadvertent loss of time services to the network.

Inventory of time dependent systems can be a little tricky. Aside from obvious systems like a primary domain controller or email server, over time it can be easy to lose track of the systems requesting the time stamps from a time server IP address. To aid in keeping track of NTP clients, the SyncServer S600 maintains a most recently used list (MRU) of IP addresses for the last 1000 unique clients. For smaller networks you can identify the IP address of every NTP client, for large networks you can identify the most active NTP time clients that might be requesting the time too frequently.

Detecting time anomalies is one of the many advantages of the Microchip Stratum 1 SyncServer. Alarms/emails/SNMP traps can be sent for everything from detecting time errors of microseconds, GPS related errors, system faults, etc. This is where exclusive reliance on time from the internet can fail your time keeping reliability; yet is also an area where time from the internet can prove useful when properly deployed. The way NTP operates through its peering/selection/clustering algorithms, setting up time associations within the stratum 1 server with a handful of internet time servers can operate as a time crosscheck of a Stratum 1 time server clock. If the local preferred hardware reference clock inside the stratum 1 time sever goes astray, the internal NTP daemon which is monitoring all clocks will detect the time error and disqualify the local clock. It will then automatically use the next best internet clock for time. If this occurs there will be a flurry of alarms from the SyncServer, but it will continue operating, just at a lower stratum, and continue to provide time services to the network.

Holdover time refers to how long a stratum 1 time server can maintain the desired time accuracy to UTC even though it has stopped synchronizing to a primary time source, like GPS for example. At the heart of this is an electronic device called an oscillator. In a perfect clock the oscillator has a perfectly stable frequency that never drifts and never allows time errors to accumulate. Those clocks are very expensive and generally owned and operated by the likes of USNO and NIST. For the typical network time server hardware clock, the oscillator is an inexpensive variety that can drift rapidly if not being routinely corrected by continuous lock to GPS. The time error drift will be even worse if it is a virtual machine.

In Stratum 1 servers, the GPS receiver is used to keep a generally better than average oscillator on frequency (called disciplining the oscillator). This way the Stratum 1 clock gets atomic clock accuracy while tracking GPS. The problem is when the GPS signal becomes unavailable for any reason, and the time server goes into holdover on the now free running, undisciplined internal oscillator. As the oscillator is now drifting it accumulates time error relative to UTC.

Reliable Stratum 1 time servers are sold with oscillator choices to accommodate the network timing requirements should GPS become unavailable. Typical offerings include the standard quartz-based oscillator, an OCXO or a Rubidium atomic oscillator. The oscillator choice is basically an insurance policy to buy time (and time accuracy) to fix a problem. The better the oscillator the less drift and the more time to resolve a problem before time accuracy of the time server becomes unacceptable for network or business operations.

A recent example of holdover compliance are electronic stock trading firms in Europe. By law they must maintain better than 100 microseconds timestamp accuracy to UTC for transaction auditing purposes. Microchip Stratum 1 SyncServers used in this application are always equipped with Rubidium atomic oscillators. During normal operations, the SyncServer is accurate to 15 nanoseconds, however, if GPS is lost, the SyncServer goes into holdover on the Rubidium oscillator. Rubidium's are very stable oscillators and the SyncServer will be accurate to better than 100 microseconds for over two weeks. This allows the IT team time to resolve a GPS related issue with no interruption in their business of stock trading or compliance with the law.

For a relative comparison of typical accumulated time error during holdover, see the chart below. You can see that a relatively inexpensive investment in an OCXO or Rubidium oscillator upgrade makes a substantial difference in time accuracy during holdover, should the time server go into holdover. When an upgraded oscillator is deployed, most NTP clients will be unaffected if the Stratum 1 time server goes into holdover, even for extended periods of time.

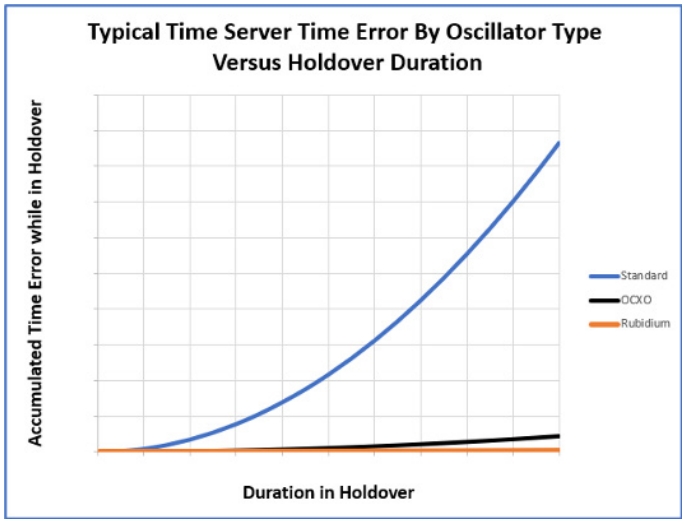

Figure 3. Typical time server accumulated time error, by installed oscillator type, versus holdover duration.

Figure 4. Microchip Rubidium based Miniature Atomic Clocks used in SyncServers for extended holdover can fit in the palm of your hand.

Degraded timing notifications are an inherent part of the NTP protocol between time servers and time clients but lack actionable notifications for IT administrators. To start, built into NTP are metrics and notifications that time servers provide to time clients that assist the client in accepting or rejecting the time from the server. What is missing is when the time is rejected by the client there is no notification, such as an SNMP trap, to the network operators. This is a fundamental vulnerability in getting time from internet-based time servers. The internet-based time server time quality may be degrading, it may be free running, its stratum may be changing, etc. but the only options available to the time client are to accept the time, reject the time, or switch to another time server that it has been made aware of. All of this will happen with no notification to the network administrators or other applications. This presents both reliability and accuracy vulnerabilities.

Real-time actionable notifications are an integral part of a SyncServer S600 Stratum 1 network time server. A long list of internal parameters regarding the state of the time, time source, GPS operations, time quality, holdover, etc. are all actively monitored by the SyncServer. Events can trigger SNMP traps, email notifications and more to notify the network operators of an issue if one occurs. This is essential when the SyncServer is operating as the central source of time for network operations and provides the means to stay ahead of network wide timing degradation.

2. Know Your Timing Source(s)

"*CISA recommends using at least two or more traceable time sources with the lowest possible stratum to minimize or eliminate timing errors*[2]*.*"

Primary and secondary sources of time are essential for reliable network time service operations. With a Stratum 1 SyncServer S600 the primary source of time for enterprise networks would be GPS/GNSS, as discussed above. GPS provides the most accurate, reliable, and secure source of time. The prudent IT administrator then configures the SyncServer to "check" the NIST internet time servers by creating NTP associations to them. This is simply done by adding the IP addresses of the publicly available NIST time servers[3] to the NTP Associations list in the SyncServer. By doing this

[2] TIME GUIDANCE for Network Operators, Chief Information Officers, and Chief Information Security Officers, page 6, CISA June 2020
[3] https://tf.nist.gov/tf-cgi/servers.cgi

the SyncServer identifies GPS as the best source of time, but also monitors, calibrates and "crosschecks" to the NIST servers. If GPS ever becomes unavailable or grossly incorrect compared to the NIST time servers, the SyncServer will switch to the next best time server automatically and send traps and notifications to the network operator. Recommended best practice is to add at least 3 NIST time servers from different locations to the NTP associations. This allows for ongoing time characterization and prioritization of those time sources.

The level of accuracy of the SyncServer S600 while tracking GPS is 15 nanoseconds to UTC(USNO). This translates into 5 microsecond time stamp accuracy to UTC for regular NTP operations, and around 15 nanosecond time stamp accuracy for the secure NTP Reflector operations. While this may be more accuracy than is required in the moment, as networked devices and the network itself gets faster over time, achievable accuracy at the time client improves due to reduction in fixed asymmetric time packet delays. Having a very accurate Stratum 1 time server to start will futureproof timekeeping across the network as various devices and infrastructure components are upgraded over time. This will be reflected in log file time stamp accuracy, SIEM report integrity, etc.

A time source with adequate holdover to prevent business and network operations disruption due to unacceptable time error is directly related to the installed oscillator in the time sever. As mentioned above, different oscillators provide different amounts of time accuracy protection versus duration in holdover. The SyncServer S600 web GUI allows the user to predefine either the acceptable time error or how long to stay in holdover. Once a value is decided, for example the desired worst-case time error, the estimated days in holdover before this time error is reached is set. This is very useful as it provides an idea of how much time IT operations will have to resolve a problem before the SyncServer reaches that user set level of time error should the SyncServer go into holdover. It is important to note that the SyncServer will go into holdover if GPS becomes unavailable, and how long it stays in holdover is user defined as a function of estimated time accuracy. Once the holdover duration expires and if there is still no GPS signal, the SyncServer will begin to synchronize with the most accurate of the user configured NTP servers. It will also drop from Stratum 1 to the stratum of the remote time server + 1.

Obtaining authenticated time from internet-based time servers is rare and requires the cooperation of the owner of the time server. A clear advantage to having a SyncServer inside the firewall is that you can use the time authentication methods available as well as a long list of other security choices, such as access control lists, multiport time operations, out-of-band management, remote user authentications, x.509 certificates, etc. All of these contribute to the overall security of the SyncServer itself and the time stamps provided.

3. Know Your Users

Regulatory compliance is increasingly including network timing operations in the areas of time accuracy and security. In stock trading FINRA requires 50 millisecond timestamp accuracy to UTC and in Europe MIFID II requires stock traders and exchanges to be within 100 microseconds to UTC. With respect to security, organizations that process large numbers of credit card transactions must comply with PCI-DSS[4] which includes time keeping standards and related security requirements.

4. Regularly Update Your System

Timing and synchronization devices should absolutely be kept up to date with respect to installed software and ongoing security hardening. If for no other reason than to assure there are no vulnerabilities related to published Common Vulnerabilities and Exposures (CVEs). By default, Microchip SyncServers periodically check to see if there is a more current version of software available than the version installed. If there is, the SyncServer will notify the network operator by way of an email, SNMP trap, etc.

Default security settings in the SyncServer are further enhanced when the network operator deploys best practices of strong passwords with expiration cycles, access control lists, remote authentication such as TACACS+ RADIUS or LDAP, x.509 certificates and out-of-band management from LAN ports providing timing services to the network.

Adaptable and upgradeable software and features are central to the hardware and software architecture of the SyncServer. All SyncServers ship with extra hardware features that are inactive until enabled through optional software licenses. Examples include adding PTP operations with hardware time stamping and expanding GNSS constellation tracking from GPS (standard) to also include the Galileo, GLONASS, BeiDou, and QZSS constellations.

5. Test System and Sources

Testing to known events, such as a forthcoming UTC leap second, may not be possible or practical in a deployed time server. SyncServers are designed and tested to work with leap second adjustments that may or may not occur in the future. When leap second adjustments are announced by the International Earth Rotation and Reference Systems Service[5], Microchip use GPS simulators to confirm the SyncServer will operate properly through the leap event and publishes the test results. Daylight savings time adjustments are not material to a legitimate Stratum 1 time server as all timing keeping is relative to UTC by definition in the NTP protocol. Local time adjustments are always handled through the operating system of the time client.

[4] Payment Card Industry - Data Security Standard; https://www.pcisecuritystandards.org/
[5] www.iers.org

6. Timing Diversification

Multiple timing sources being referenced is easily accomplished by using GPS/GNSS as the primary, providing local Stratum 1 caliber timing, and as mentioned in Section 2, using time from the NIST internet accessible time servers as a backup. This configuration in combination with an upgraded oscillator as reviewed in the previous discussion on holdover, results in a very robust solution to provide accurate and reliable timing services to ensure ongoing network and business operational integrity.

7. Anomaly Detection

Anomaly detection in a time source is easiest done when that time source is a dedicated Stratum 1 server inside the firewall and well equipped to provide notifications, alarms and associated log files for any anomaly that might be detected. Furthermore, this is not possible with time sources from the internet as they are merely an IP address that returns a time stamped NTP packet on request. Beyond that, you really know nothing about it.

Corrupted or unavailable time sources have been accommodated since the advent of NTP itself. In the event that access to the time from a Stratum 1 server is not available, such as a disconnected network cable, the guidance is twofold: First, have more than one time source at all times, preferably two Stratum 1 time servers; and second, make sure essential time clients have access to both. This way in the event of the loss of one time source the time clients will automatically switch to synchronizing with the other. This how NTP operates as it adds resiliency to network timing operations and is a longstanding best practice for network operators.

Example Timing Topology

Using the diagram from earlier in this document, the following is a recommended deployment scenario.

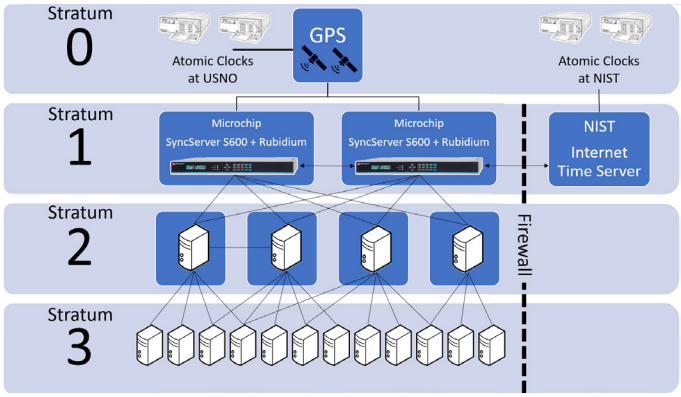

Figure 5. A typical enterprise deployment of SyncServer S600 Stratum 1 network time servers with atomic clocks synchronizing to GPS inside the firewall and using Internet NIST time servers as a fallback timing source.

The Stratum 1 SyncServer S600 network time server with a Rubidium atomic oscillator for holdover will provide accurate, secure, and reliable time to the network. With two SyncServers there is redundancy in that if one is denied a GPS signal it will synchronize over the network with the other. If GPS is lost altogether the SyncServers will go into holdover on the internal Rubidium oscillator and very accurately keep the time while the GPS issue is resolved. If the holdover duration expires the SyncServer will revert to getting time from the best NIST internet time server that it has been time qualifying continuously in the background. All the while the entire network will be kept synchronized. The Stratum 2 layer and below devices are configured to synchronize with more than one time server in the event there is a complete loss of an upper layer clock on the network. This implementation can also be structured with the SyncServers located in different geographies with both synchronizing to GPS but still connected to each other over the network. This is a common enterprise deployment for multi-campus time distribution.

Summary

Accuracy, Reliability and Security are the most commonly cited reasons to install a Stratum 1 SyncServer S600 network time server. They are affordable, easy to install, extremely accurate and secure, and unfortunately easy to forget that you own due to their high reliability. Which brings you full circle back to the CISA Time Guidance publication which correctly emphasizes regular review of the timing infrastructure on your network from the timing source down to the applications that rely on accurate timekeeping for your business.

DS00003612A