

The Importance of Network Time Synchronization for Enterprise Networks

As you read this, your network of workstations and servers, each with its own clock, are actively timestamping files, emails, transactions, and so on. At the same time, your server logs are recording every type of transaction in case the information is needed for audit or forensic analysis. At some point during the day, it is quite likely that automatic processes (such as archiving, directory synchronization, and cron jobs) will execute and alter files based on these timestamps.

Fundamental to all this is the belief that the time is correct. But is it? To answer that question, you must first consider these questions:

- Is your time source accurate and precise?
- Is your time source secure? That is, is the time itself vulnerable or is the time source vulnerable?
- Is your time source reliable?

In many of today's networks, the network engineering answer to these questions is often, "I think so." In truth, computer clocks are notorious for drifting. They are typically based on inexpensive oscillator circuits or battery backed quartz crystals that can easily drift seconds and minutes per day, accumulating significant errors over time.

Many organizations get "free time off the Internet" which, while free, poses serious risks in terms of time accuracy, reliability and network security. After all, it is not your clock, you have no idea if it is 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 know very little about an internet time server.

Your Network Probably Needs Reliable Network Timing

Quality time in an enterprise network is critical to operate the network in a reliable and secure manner as well as to support key applications. Compliance and forensics are also key drivers to ensure optimal network timing performance. Often it is not until a network fault occurs that organizations become painfully aware of the importance of time synchronization to diagnose the problem.

Table 1 lists the key areas where quality network timekeeping is required in both network operations and applications. If any of these applications are important on your network, then network timing is also important on your network and you should read on.

Network Operations & Applications That Rely on Accurate Time	
Log file accuracy, auditing, and monitoring	Transaction processing
Network fault diagnosis and recovery	Email
Virtual environments	Legal and regulatory requirements
Directory services	Scheduled operations
Access security and authentication	Real-world time values

Table 1: Key areas requiring accurate and secure network timekeeping.

Log File Time Accuracy, Auditing, and Monitoring



Server log files and subsequent reports enable assessment of network activities. This includes firewall related activity, bandwidth usage, trending, various logging, management, authentication, authorization, and accounting functions. Because server logs are a compilation of information from different hosts, accurate timestamps are essential for ordering events, and identifying and troubleshooting root-causes to network issues.

Speed Network Fault Diagnosis and Recovery



Most IT organizations are measured on their ability to maintain full flow network operations. In the event of a failure, accurate network timing is crucial for fault diagnosis and recovery. To assist in fault diagnosis key network events are trapped, reported, and logged (typically using syslog services that accumulate messages from servers, routers, etc.). Should the network fault, a root-cause analysis is initiated, looking through the reported stream of time-stamped events. If these timestamps are synchronized, the proper order can be established, and the root cause quickly identified. Root-cause isolation can be obscured and downtime prolonged without accurate network time synchronization to provide the accurate timestamps.

Resolve Virtual Environment Timekeeping Challenges



Within virtual environments such as VMware and Hyper-V, guest virtual machines (VMs) share a host's physical resources, which affect the guests' ability to keep their own accurate time. Accurate time in VMs depends on regular servicing of timers so that the VM clock continually moves forward smoothly. Unfortunately, a VM must often wait while the host OS is busy servicing other guests, typically resulting in the VMs' time falling behind or exhibiting erratic time behavior due to missed clock ticks. As with a real machine, the integrity of all operations of a virtual machine depends on keeping the time up to date (in other words, on not relying exclusively on the host resources). A dedicated NTP time server coupled with good NTP clients works to offset the negative timekeeping effects prevalent in virtual environments. While the guest VM is subjected to irregular clock ticks from the host, good NTP client software running on the guest can compensate for and work to overcome those timing irregularities.

Directory Services Need Time Sync



Many network directory services systems exchange information and synchronize changes in the directory services database according to timestamps. Without a time-synchronized network, time-sensitive systems and applications will not work correctly. In a Windows active directory network for example, all Primary Domain Controller (PDCs) and client workstations need to synchronize with a single, accurate, and standard time source.

Access Security and Authentication Need Sync



Synchronized time is critical in Windows because the default authentication protocol (Kerberos) uses workstation time as part of the authentication ticket generation process. Windows includes the W32Time Service tool to ensure that all Windows-based computers in an organization use a common time. The time service uses a hierarchical relationship that controls authority and does not permit loops so as to ensure appropriate common time usage. This continues through the hierarchy of domains to the PDC at the root of the tree. This PDC is set to synchronize with a reliable time source, such as a dedicated network time server. If a time server is not available and the resulting time difference between domain controllers drifts beyond the skew allowed by Kerberos, authentication/logon between domain controllers and clients may not succeed and thereby impair the proper operation of the network.

Time Coordinated Scheduled Operations



Cron scripts and crontabs are a list of one or more commands to a computer operating system or application server that are to be executed at a specified time. These commands are commonly data backup oriented and happen at prespecified times. Synchronization to a single host with an acceptable time source is mandatory so that commands run when expected. In the case of multiple hosts responsible for executing independent cron files, time synchronization between the hosts becomes even more critical to ensure that scheduled activities are properly coordinated.

Real-World Time Values



There is no substitute for operating a network using real-world time values. While you can synchronize a network to the incorrect time and make it work, this is a very undesirable policy. Local networks are typically interconnected with other larger networks and correct time is a single common denominator. Real-world network time is based on Universal Time Coordinated (UTC). Networks operating on the underlying UTC share a common time base. UTC time is best obtained from an accurate, secure, and reliable source, and then converted to local time by operating systems that reference that source. It is this common time reference that provides network operators the meaningful time-accurate information they need about their network.

Accurate Transaction Timestamps



Time synchronization in business transaction processing is essential, especially when processing is distributed among cooperating systems on a geographically distributed network. The need for millisecond or microsecond timestamp accuracy derives from the need to execute or report transactions in a correct sequence, particularly if many transactions occur almost simultaneously, resulting in a need for sub-millisecond resolution. And, because these systems can be geographically separated by great distances, keeping accurate time to UTC becomes mandatory.

Email Timestamp Integrity



Email is the de facto standard of written business communication. Every email message that passes across the network bears the originator's timestamp. If that timestamp is incorrect, it can create confusion on the part of the recipient and challenge the credibility of the originating organization.

Legal and Regulatory Time Accuracy Requirements



Accurate and traceable network time can be found in industry regulations. In the United States, for instance, the Financial Authority Regulatory Authority (FINRA) that governs stock trading requires its members to timestamp stock trades with an accuracy of 50 milliseconds or better traceable to UTC at the National Institute of Standards and Technology (NIST). In Europe, the MiFID II legislation is an even more stringent 100 microseconds to UTC. The reason for synchronized, traceable time in this application is to validate when a stock transaction occurred for the purposes of order auditing. Other industries like the payment card industry (Visa, Mastercard, etc.) also have time related requirements as part of their network data security standards.

The Answer: A Stratum 1 Network Time Server Inside Your Firewall




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

Accuracy and security is why many enterprise networks today rely on Stratum 1 network time servers located inside the firewall that acquire time from Global Navigation Satellite Systems (GNSS) and distribute it to clients over the network through the Network Time Protocol (NTP).

For reliability, these NTP servers also employ stable internal oscillators in case the server loses the satellite signal and goes into holdover. Oscillator holdover accuracy varies based on the type of oscillator: from 400 microseconds over the first 24 hours of signal loss for a standard quartz oscillator to less than 1 microsecond for an inexpensive rubidium atomic oscillator.



Figure 2: Microchip Rubidium based Miniature Atomic Clock used in Stratum 1 SyncServers for extended holdover can fit in the palm of your hand.

Accuracy, Reliability and Security



The winning trifecta of enterprise network timing for business and network operations is accuracy, security and reliability. These are the most commonly cited reasons to install a Stratum 1 network time server, such as a Microchip SyncServer S600, inside the firewall. SyncServers are affordable, easy to install, extremely accurate and secure, and so reliable you'll likely forget you own it.