



Does Anybody Really Know What Time it is?

Copyright © Brilliant Telecom, January, 2007
Reprinted June 2008, used with permission

*As I was walking down the street one day
A man came up to me and asked me what the time was that was on my watch, yeah
And I said*

*Does anybody really know what time it is
Does anybody really care
If so I can't imagine why
We've all got time enough to cry*

*And I was walking down the street one day
A pretty lady looked at me and said her diamond watch had stopped cold dead
And I said*

*Does anybody really know what time it is
Does anybody really care*

Chicago, 1970 – Does Anybody Really Know What Time It Is

While Robert Lamm may have expressed his sentiment about time in 1970 with his hit “*Does Anybody Really Know What Time It Is*”, the rest of the world does indeed care about time. From your alarm clock going off far too early each morning to your paycheck making it to the bank when you expect it to, or even ensuring you’re ready for the opening kickoff of Monday Night Football, we are all concerned about time.

These may be simple everyday examples, but the point is easily made – Time is Important. It’s not enough to simply know what time it is, there are a multitude of policies and regulations that mandate accurate time and timestamps. The New York Stock Exchange (NYSE) mandates the following:

"New Rule 132A requires members to synchronize the business clocks used to record the date and time of any event that the Exchange requires to be recorded. The Exchange will require that the date and time of orders in Exchange-listed securities to be recorded. The Rule also requires that members maintain the synchronization of this equipment in conformity with procedures prescribed by the Exchange".

Specific NYSE Time Synchronization Requirements

Rule 132A contains two specific requirements:

- ***Clocks Synchronized to Commonly Used Time Standard***
All computer clocks and mechanical time-stamping devices must be synchronized to a commonly used time standard, either the National Institute of Standards and Technology (NIST) or United States Naval Observatory (USNO) atomic clocks.
- ***Synchronization must be maintained***
Rule 132A also indicates that the member must ensure that their systems remain synchronized.

Other business and government regulations have similar requirements for accurate timing and accountability. In fact, not only is the time stamping capability vital, the auditing and traceability of those time-stamps is also critical. Recent US regulations such as the Sarbanes Oxley (SOX) and the Health Insurance Portability and Accountability Act (HIPAA) mandate both an accurate time-stamp for all transactions and documents, as well as an audit trail of such transactions. Organizations that fail to meet these requirements face stiff fines and even jail time for individuals responsible for those failures.

Regulations, fines, and moldy jail cells aside, there are other reasons for accurate time. What about manufacturing and chemical processes? On an assembly line, ensuring that parts arrive exactly when they are needed is critical, particularly as they become more and more heavily automated. If a robotic

arm reaches for a specific rivet to attach a body panel to the frame of a car and it isn't there, the integrity of the automobile is compromised.


In the case of using time to determine location, such with the Global Positioning System (GPS), a discrepancy of just one second equates to an error in position of .23 nautical miles at the equator.

Accurate Time

Accurate time; how is it assured -how do you know your alarm clock is accurate? Mechanical clocks are notoriously unreliable, but most electronic clocks also keep rather inaccurate time. One reason for this is that designing a computer or electronic device to keep accurate time is not often a core competency of the device or computer manufacturer, nor is it primary function of that device.

However, even reasonably accurate computer clocks vary due to manufacturing defects, changes in temperature, electric and magnetic interference, the age of the oscillator, or even computer or electrical load. In addition, even the smallest errors can add up over a long period to a significant difference in time. For example, two clocks, synchronized at the beginning of the year, but one running only small amount slower -say an extra microsecond to increment a second; by the end of a year, the two clocks will differ in time by more than 30 seconds. If a clock is off by just 10 parts per million, it will gain or lose almost a second a day. These measures are actually fairly optimistic examples of the accuracy of some of the clocks in modern workstations and PCs and certainly better than the alarm clock on your nightstand. Luckily, we've come a long way and are no longer reliant on 17th century time keeping devices; in fact, time, and it's adjunct, frequency are the most accurate measurements we can make. The measurement of the duration of a second is so precise it has become the standard by which other units, such the volt, ampere, ohm, and meter are defined.

The ultimate in modern time references are capable of maintaining timing as accurate as ± 1 second in 10 million years.



Unless you have the luxury of your own atomic clocks throughout your network, having any sort of meaningful time synchronization is almost impossible if clocks are allowed to be free running. In some environments, like your nightstand, this lack of synchronization isn't a big issue. However, in most modern networked computing environments, time synchronization is critical. To reduce confusion in shared networks, it is crucial for the modification times to be consistent, regardless of what machine the file-systems are on. Billing services and similar applications must know the time accurately. Sorting email and other network communications can also be difficult if time stamps are incorrect. In addition, tracking security breaches, network usage, or problems affecting a large number of components can be nearly impossible if time stamps in logs are inaccurate.

Applications such as cryptographic key management and secure document transmission require using accurate, as well as video transmission encoded time stamps which match un-encoded time stamps to help assure document authenticity. For instance, secure Remote Procedure Calls (RPC) needs clocks to be synced to within 15 seconds for proper operation. In addition, interactions with dynamic events such as TV programming require careful synchronization of time.

Accurate timing devices, like your own personal atomic clock, can be purchased, but for the vast majority of organizations, there is a simpler and more cost effective method. The most commonly used method of ensuring accurate time; at least where there is a computer network, is the Network Time Protocol, or NTP.

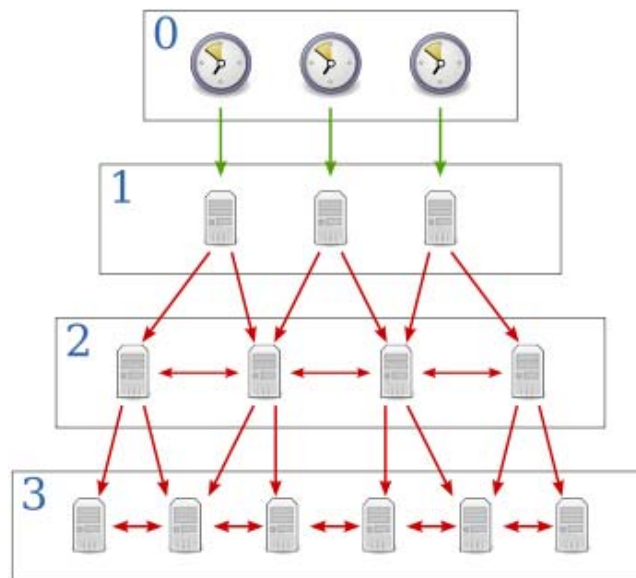
Network Time Protocol – RFC 1305

NTP is not based on the principles of synchronizing machines with each other. Instead, it is based on the principles of having all machines get as close as possible to the agreed upon reference time.

At the top of any NTP hierarchy are one or more reference clocks. These are electronic clocks synchronized to a common time reference and each other using some methods outside the scope of NTP, such as the atomic clocks mentioned above. These reference clocks are assumed to be accurate – and the methodology they use to maintain synchronization with each other is beyond the scope of this document; for additional information please refer directly to the RFC, which may be found at <http://rfc.net/rfc1305.html>.

The accuracy of all other clocks is judged according to how “close” a clock is to a reference clock (the Stratum of the clock, as described below), the network latency to the clock, and the claimed accuracy of the clock.

NTP works on a hierarchical model in which a small number of servers give time to a large number of clients. The clients on each level, or Stratum, are in turn, potential servers to an even larger number of clients of a higher numbered Stratum. Stratum numbers increase from the reference clocks (Stratum 0) to the low numbered strata at the leaves of the tree. Clients can use time information from multiple servers to automatically determine the best source of time and prevent bad time sources from corrupting their own time. Figure 1 illustrates the hierarchical strata model of servers used in NTP.



Servers that are directly connected to a reference clock are termed Stratum 1 servers. A reference clock connected to a Stratum 1 server is referred to as a Stratum 0 server. Clients never communicate directly with a Stratum 0 server; they always go through a Stratum 1 server synchronized to a Stratum 0 server. Regardless of the source of time for a server, it is important to remember that the accuracy of these time signals can vary widely. Stratum 0 clock references coordinate time and compare drift and skew amongst each other to maintain accuracy.

For the rest of us, there are several public NTP servers available on the Internet. They use Coordinated

Universal Time as their ultimate source of time. UTC evolved from Greenwich Mean Time (GMT), and still uses the Greenwich time zone as the zero offset. GMT, which is based on the earth's rotation, is not constant enough to be used for detailed time measurements (the earth is currently spinning about 1 second slow – it's not known when it will catch up). UTC is based on a standard second length determined by quantum phenomena (e.g. the radio-active decay of Cesium atoms, ala the Royal Observatory).

In environments that need more accurate time than an Internet link will allow (due to latency, service provider restrictions, or other concerns), or environments that cannot rely on Internet time sources due to security implications, something else is clearly required. Synchronizing a few machines to an arbitrary time source, such as the internal clock on a given server, may be acceptable in a few rare cases, but in any sort of large installation it is critical to keep the clocks synchronized with some maintained time standard. Regardless of the configuration, an NTP server needs to be set up in order for clients to use it for synchronization.

NTP Clients and Servers

The relationship between NTP servers and clients can be configured to operate in several ways. Computers using NTP can operate in different modes with respect to different machines. For instance, a single machine could be a client of a machine with a lower Stratum number, while being a peer to a machine on the same Stratum, and a broadcast server to a number of clients at a higher Stratum number.

Server – An NTP server serves time to clients. Clients send a request to the server and the server sends back a time stamped response, along with information such as its accuracy and Stratum.

Client – An NTP client gets time responses from an NTP server or servers, and uses the information to calibrate its clock. This consists of the client determining how far its clock is off and adjusting its time to match that of the server. The maximum error is determined based on the round-trip time for the packet to be received.

Peer – An NTP peer is a member of a group of NTP servers that are tightly coupled. In a group of two peers, at any given time, the most accurate peer is acting as a server and the other peers are acting as clients. The result is that peer groups will have closely synchronized times without requiring a single server to be specified.

Broadcast/multicast server – An NTP server can also operate in a broadcast or multicast mode. Both work similarly; broadcast servers send periodic time updates to a broadcast address, while multicast servers send periodic updates to a multicast address. Using broadcast packets can greatly reduce the NTP traffic on a network, especially for a network with many NTP clients.

Broadcast/multicast client – An NTP broadcast or multicast client listens for NTP packets on a broadcast or multicast address. When the first packet is received, it attempts to quantify the delay to the server in order to better quantify the correct time from later broadcasts. This is accomplished by a series of brief interchanges where the client and server act as a regular (non-broadcast) NTP client and server. Once these interchanges occur, the client has an idea of the network delay and thereafter can estimate the time based only on broadcast packets.

Threats to Accurate Time

The notion of accurate time is essential to determining the order in which events occur. This is a fundamental aspect of transactional integrity, system and network-wide logging, auditing, troubleshooting and forensics. Having an accurate time source plays a critical role in tracing and debugging problems that occur on different platforms across a network. Events must be correlated with each other regardless of where they were generated.

Furthermore, the notion of time (or time ranges) is used in many forms of access control, authentication, and encryption. In some cases, these controls can be bypassed or rendered inoperative if the time source could be manipulated. For example, a payroll function could be tricked into providing access over a weekend when normally it would be restricted to normal business hours.

Many organizations have become reliant on NTP just as they are with other services such as the domain name service (DNS). This reliance can be a weakness if the service is not properly safeguarded. Therefore, it is important that these time sources are adequately protected against a wide array of threats, internal and external, local and remote. Time is not just an extraneous service. It is fundamental to the successful operation of today's environments.

The most significant risks to NTP services are tampering, jamming and Denial of Service (DOS) attacks. Tampering occurs when the NTP server is affected by either accidental or malicious data modification. Jamming occurs when a time server is either destroyed or prevented from providing NTP service. DOS attacks occur when an NTP server is flooded with traffic, either NTP requests or other (e.g. management) traffic, and is unable to validate NTP requests. As with any other application, administrators must remember that NTP is not guaranteed to be secure; poor coding and other flaws in the program could allow unintended access to NTP internals or the underlying operating system. NTP servers may be capable of protecting themselves against some of these threats using architectural choices such as redundancy, and configuration options such as access control and authentication.

Service Providers and NTP

Nearly all network devices, PCs, printers, hubs, routers, etc., generate NTP requests. Most request time from a handful of public domain NTP servers, or in the case of Windows PCs, at time.microsoft.com. As you might imagine, with billions of network devices all requesting time, this adds up to a significant amount of traffic. In extreme cases, some consumer devices are known to request time updates every 2 seconds.

Many service providers view these NTP packets (UDP port 123) as unnecessary noise traversing their network and either block them entirely or establish proxies at the edge. An NTP proxy would be a process running on the router that your home or enterprise network connects to; any NTP requests generated by your network would be intercepted and responded to at that point. All is well and good, assuming you trust your ISP to have configured everything properly throughout their network and to have provided traceability back to the ultimate reference clocks.

Unfortunately, that is not commonly the case. Your NTP request is responded to locally, typically by a router whose own clock was set once, the last time it was rebooted, and is now free-running. Or, just as bad, pointed at another router that is in turn pointed back. What this means is that even though you have NTP running, you are in actuality, no better off than having a good clock in a free-running state.

What is needed is a local, secure, highly accurate and traceable timing source. In other words, a GPS based NTP server that is controlled by the local enterprise. Such devices exist – such as the CXR Larus line of StarSync™ and StarClock™ products.

For further information or to discuss how we can help solve your timing problems, contact us at www.cxrlarus.com, or phone us at 408-573-2700.

By the way, does anyone have the time?

"Now, perhaps more than ever, time is fundamental. In today's fast-paced environment, we have placed a premium on efficient time management -- meetings, schedules, dates, being on time, taking time, giving time, setting a clock's time, and on and on. To that end, we have developed incredibly accurate clocks to help us turn time into an extremely useful tool."

GPS World, November 1998

