September 22, 2004

Well It's About Time

Vance Rider
GIAC Security Essentials Certification (GSEC)
Practical, version 1.4c, option 1

Abstract

This document describes the role time plays in a networked environment. Specifically it introduces the reader to the Network Time Protocol (NTP) and how it is used to synchronize computer clocks together via a hierarchical master slave relationship in a secure manner. We peer into the definition of standard time and who some of the purveyors of trusted time are. A brief overview of how a typical Windows environment keeps up with an external time source is examined. Upon conclusion, we take a closer look at how time is disseminated to networks all around the world and the various hardware devices involved.

In The Beginning

Network Time Protocol (NTP) was developed by David L. Mills PHD at the University of Delaware. It was in its infancy stages as early as 1981[1] but made its official debut with version 1 in 1988[2]. NTP version 3 is the currently accepted standard as described in [RFC1305][3] The evolutionary step to the version 4 [RFC 2030][4], has been made and is now being deployed.

Time for an Upgrade

You might be wondering…if NTP v3 has been out since the early 1990's isn't it a bit outdated? Well, the answer to that is yes and no. NTP v3 has enjoyed a reputation as being fairly solid however version 4 started being developed in the early 90's as well. It has just taken a painfully long time to add the extra features and make requested modifications. With the many evolutions of just the fourth version, convention will allow us to refer to it using the point scale upgrade .x.x i.e. NTP v4.x.x . Systems administrators have been using NTP v4.x.x with greater frequency since 1998-99. Until it officially steps out of the development range it will not be fully maintained by the team of experts charged with authoring NTP v4.x.x.

As you might imagine, one of the driving issues for enhancing NTP v3 is security. Version 3 does use cryptographical means to ensure it is communicating with a trusted source. However, it achieves this through symmetric-key[5] cryptography. The primary drawback to this process is the requirement for both parties to use a shared secret key. This key must then be distributed to each party in advance. This can be subject to malicious activity especially when the key is used by a large number of clients.

NTP v4.x.x has greatly improved on this security model by including support for both symmetric and public-key[6] (Autokey)[7] cryptography. The chief advantage of using a public key is that it provides an added layer of defense in depth, thereby increasing the security against an intruder or imposter. This is a vast improvement but it came at a price that not everyone was prepared to dole out. In order to verify and reliably construct server identification credentials along with the public certificates means consuming additional available processor recourses. This has been considerably improved by revamping the specific algorithms that previously taxed the processor resources.

Please note the following OpenSSL excerpt from the NTP Version 4 Release notes:

> As required by Defense Trade Regulations (DTR), the cryptographic routines supporting the Data Encryption Standard (DES) have been

> removed from the base distribution of NTPv3. For NTPv4 a new interface has been implemented for the OpenSSL cryptographic library, which is widely available on the web at www.openssl.org. This library replaces the library formerly available from RSA Laboratories. Besides being somewhat faster and more widely available, the OpenSSL library supports many additional cryptographic algorithms, which are now selectable at run time. Directions for using OpenSSL are in the Building and Installing the Distribution page[8].

You should definitely be using NTP v4x.x over NTP v3 if your circumstances allow such a choice.

Besides better security you will also find other important features. Below are just a few of the more important ones.

- The new addressing schema, IPv6 is fully supported as well as the current IPv4.
- The way calculations were done in the past has made advancements; no longer will you find the 64-bit fixed point format, rather the 64-bit floating double format to help reduce algorithm size and improve speed.
- The way clock discipline algorithms were executed have been retrofitted to improve accuracy and reduce network jitter.

Back to True Time

Mills recognized the need for computer systems to know 'true time' - a timestamp recognized as a single and authoritative value to be shared between two or more systems. The paradigm he chose was hierarchical with the reference source at the top providing true time to lower nodes which would then act as an authentic time source to the next lower node and so on.

Each tier is then referred to as a stratum. Simply stated, stratum is defined as how close a clock is to a reference clock. The achieved effect is that you have a small number of servers providing accurate time to a larger number of clients. In this model the reference clock would be dubbed stratum 0 and next lower clock named stratum 1, the next stratum 2 etc.

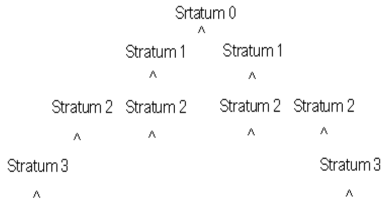

The maximum NTP stratum number for a client is 15; however in practice it is rare to find clients with a stratum above 4 or 5 in most real-world configurations[9]. Conventionally, lower level stratums should be required to synchronize with their parents, never directly with a stratum 0 server.

Stratum 1 clocks usually synchronize to some national time source by way of satellite, radio or modem. (Discussed in detail in later chapters.) If your network requires a very high degree of security and/or autonomy then you may consider purchasing your own time server to provide your own stratum 0 services. Such hardware devices are available from vendors but are very expensive.

Typically, one would use a stratum 1 server made available to the public[10]. Because many of these services are generously provided by private companies synchronization should be reserved for stratum 1 and 2 server update requests from your network, not your clients - as this would impose greater use of their resources.

When setting up your NTP server it is advisable to link to multiple well synced stratum 1 or 2 servers to ensure sustained accuracy. It is possible in some situations that a stratum 1 server is simply using its own internal computer clock as a trusted source, producing a less than desirable result. Therefore, always use several verified sources.

Time

So where does this time come from you might ask? Usually when we refer to time in a global sense we use the term Greenwich Mean Time (GMT). This is understood to be the 'zero offset' on which we base the different time zones. Because GMT is calculated using the rotation of the earth we can not trust its accuracy. The reason is due to earth's rotary motion; it is not constant enough to be used for exact measurements. So, most public stratum 1 clocks available via the internet utilize Coordinated Universal Time (UTC), instead. (Did you notice the acronym is actually UTC?) UTC evolved from GMT; as such, people often believe the two to be one in same, but this is not the case.

Who keeps the time? The International Bureau of Weights and Measures is the authoritative source for UTC. They poll data from the timing laboratories to provide this trusted international standard. Using this data they are able to approximate UTC to the nanosecond, (a billionth of a second), per day. The length of a UTC second is defined in terms of an atomic transition of the element cesium under specific conditions, and is not directly related to any astronomical phenomena[11].

Where do some of these public time servers exist? The National Institute of Science and Technology (NIST) is one such source as it hosts about two dozen. Others can be found at the US Naval Observatory (USNO) and the Canadian Meteorology Centre (CMC). For an up-to-date list of publicly maintained time servers browse the Public NTP Time Servers web page[12].

It should be stressed that access policies are usually in effect for most upstream time servers so be sure to fully understand and observe the rules carefully. Also, when making your selection try to pick ones that are geographically and ‘network topology’ close to your proximity as this will help reduce latency and network traffic. Furthermore, while keeping vicinity in mind, a prudent network administrator will conceive plans to ensure redundancy is in place. This can be achieved by picking individual servers routed via different networks, thus safeguarding time sensitive and mission critical applications.

Time for an Example

NTP software has been written for almost every kind of operating system platform commonly used today. Unix, Windows, VMS are just a few. In fact many embedded systems rely on and trust NTP as well. One example is NASA, it's what the space shuttle and other spacecraft use[13].

Let's make a closer examination of one of these applications and see how NTP plays its role. In our example we will use the Microsoft Windows 2000 and Windows XP Professional operating systems located together in a single domain. The public time server we will link upstream to will be from NIST.

Microsoft uses an architecture based on the notion of a Domain. A Windows domain is used where there is a need for security and file sharing between servers and workstations. All members belong to a single domain. (There is much more to a domain but for our purposes this definition will suffice.)

As you can imagine secure network communications within a domain require a common time standard. Each computer does have its own time-of-day clock which is used to derive a timestamp and then associated with each event. But what if a certain shared event between computers in the domain had different timestamps affixed to it? Depending on which computer you were using to view the event you would observe a dissimilar time stamp.

Think of the nightmare it would be trying to forensically troubleshoot a network issue. Even if all computers were meticulously set up with the same time, say for example every six months, you would still encounter trouble as computer clocks drift.

Windows addresses this issue by providing a program that runs covertly in the background identified as a ‘service’. The Windows Time Service (WTS)[14] is responsible for synchronizing each system's computer clock with all the other systems residing in the domain. (More on that in a moment) Each time a shared file is created or changed or an electronic mail message is sent and received it would be very difficult if not impossible to be certain of the exact time the event took place with out WTS. Perhaps we have an employee time card application that is accessed via the computer or some other time-sensitive transaction.

The computer clocks must be set to use an established mutual time *and* kept in synchronization. Microsoft expects the computer clocks populating a domain to be synchronized within 2 seconds of each other.

Remember, maintaining synchronization is what NTP is all about. Typically there will be one sever considered to be the authoritative time keeper for the domain, it sits at the top of the hierarchy and usually functions as a stratum 2, and all other servers and workstations (clients) will request their time from it using WTS in *conjunction with* NTP. The authoritative server in turn is usually configured to poll several stratum 1 servers as mentioned earlier. In this case we want one of our time sources to originate outside of our network, preferably from the NIST.

NIST disseminates its time from its atomic clocks to the public using several methods. One such way is using broadcasts via short-wave and long-wave radio which can be picked up and used by your system if the proper hardware (antenna) has been installed. Perhaps more favorable to computer users are two other means provided by NIST, telephone dial-in services (ACTS) and, for those with internet access, Internet Time Service(ITS). NIST makes available a free ITS client for Windows. It can be downloaded here: http://www.boulder.nist.gov/timefreq/service/pdf/nts.pdf

Okay, so this Windows – NIST scenario is bit over simplified but you get the general idea. Without NTP there would have to be some other way synchronize the time.

Let's Get Synchronized

In general there are four ways to synchronize a network to public available stratum 1 servers: Global Positioning Satellite (GPS), Internet, Modem and Radio Frequency waves (RF).

GPS

Throughout history mankind has struggled with locating his position on the earth's surface. While many ingenious contrivances and techniques had been pioneered most proved to be futile and cumbersome. In the days of Galileo it was the heavenly stars that provided guidance to that far-off port of trade. Later in the 20th century, it was man-made stars that would that would steer men to the moon and back thus heralding in a known exact location only dreamed about in earlier times. This dawning of this new age gave birth to the Global Positioning System.

The global positioning system is an arrangement of 24 satellites networked together via radio-navigation to each other and their ground stations. They orbit the earth at a distance of about 11,625 miles. It was developed as a navigation system to establish a precise time and location by the United States Department of Defense(DOD). As such the DOD owns and maintains these satellites paid for by U.S. tax dollars.

These man-made celestial objects are used to reference each other and then calculate accurate positions on the earth to within a few meters. The underpinning of GPS is "triangulation". Very high orbits help facilitate the triangulation process as it allows a greater number of satellites to share a line of sight with each other. Using radio signals, the GPS receiver measures how far away it is from another satellite.

Including itself, four satellites are the minimum required to determine its exact position. Theoretically you could get by with three however a fourth is required because receiver clocks are not perfect. Keeping in mind that three perfect measurements can establish a certain location in a 3-dimensional space then a fourth satellite, given that location should always be able to point to that same space. Since any deviation from universal time will adversely affect all of their measurements the fourth receiver can seek out a single correction factor that it can subtract from all its timing measurements thus causing them all to again intersect at a single point. Therefore the fourth satellite provides a failsafe feature.

Global positioning satellites are broadly utilized today by many users throughout the world. As a security precaution the U.S. government restricts the precision and only allows the public's accuracy range to fall within a few meters. The DOD is actually the agency oversee GPS, and as such, it is up to its discretion how the data is disseminated. Make no mistake about it, the U.S. military is fully capable of pin-pointing measurements to less than a centimeter when necessary. However what we are really concerned about here is time, not necessarily the place, In contrast to GPS receivers, which give us an exact location, is the GPS clock whose main purpose in life is making timing data available at precise intervals and frequency.

Residing within each satellite is an atomic clock. Atomic clocks measure time by counting the oscillations of a certain type of atom, namely cesium. The clock actually measures the time by the natural vibrations of the atoms, approximately 9.2 billion "ticks" per second.

There is a new clock being developed at NIST that you should be aware of. It is called the Chip-Scale clock. Laura Ost with NIST explains how it works: Cesium vapor is confined in a sealed cell and probed with light from an equally small infrared laser, which generates two electromagnet fields. The difference in frequency of these two fields is tuned until it equals the difference between two energy levels of the atoms. The atoms then enter a "dark state" in which they stop absorbing and emitting light; this point defines the natural resonance frequency of cesium[15]. In order to receive the time from one of these satellites it is necessary to incorporate a GPS receiver into your system.

The satellite transmits a faint twenty-watt signal which is then picked up by the receiving device located at your site. Satellites broadcast a number of spread spectrum codes but only one of them, referred to as Coarse Acquisition (C\A) code is easily accessible for civilian use[16]. Once the signal has been received your computer will then further digest and calculate the velocity and position of the receiver in relation to the satellite. The best quality GPS receivers are able to measure the C/A code to better that nanosecond precision. The industry standard NMEA protocol is the protocol used by many GPS receivers in a typical network time synchronization scenario[17].

Conventional GPS receivers communicate via a standard RS-232 serial connection which can be plugged into the serial port on your computer. The connector itself can be either a 25 pin conductor DB-25 or a 9 pin connector DB-9 (also known as DE-9). Newer receivers now incorporate Universal Serial Bus (USB).

Whether you use RS-232 or USB to connect to the computer you will need to decide what kind of receiving device to employ on the other end. You have the choice of antennae: an indoor or outdoor radio receiver.

The GPS antenna has been the mainstay for quite a long time, as such it makes a good choice.

When positioning the GPS antenna on your site be sure to locate it with a decent and liberated 360 degree view to the sky. The receiver needs to be able to triangulate at least three satellites so the higher the better. Obviously the top of a building would be ideal. In practice however, the antenna will generally do just fine if attached on a window-ledge or even just near a window. Keep in mind that the signal can be substantially weakened if your window is near a lot of heavy foliage. Also, if television or radio broadcasting equipment resides close by that is not maintained properly the signal may be rendered inoperable.

You are limited in proximity to your computer by the cable length, typically about 25 meters. (Refer to your manufacture's specifications.) You can see an example of both outdoor and indoor antenna here: http://www.brgprecision.com/opgo.html http://www.brgprecision.com/opgp.html

There is one other item to consider when choosing a GPS synchronization scheme. Even though it is highly accurate your results may vary depending on both the computer and receiver. The receiver you use may indeed be able to produce a very fine outcome to within several milliseconds of UTC; but the time that is finally synchronized to your PC may not be. If your computer system does not have enough resources available or if it functions inconsistently then its ability to match the accuracy of the receiver will be limited. You should strive for a maximum of plus/minus 0.1 seconds difference between the two devices.

Radio

Radio devices have been around even before GPS. And as with GPS, they provide a good alternative for highly secure or remote sites. It is based on a simple arrangement of one trusted and authoritative source broadcasting its time, based on UTC, via the open air waves to whom ever wishes to receive it. In the United States there is one official time source responsible for transmitting to radio receivers in search of an atomic clock to sync up to. It is the NIST Radio Station WWVB, located in Fort Collins, Colorado[18]. Their signal is broadcast on a frequency of 60 kHz.

The bandwidth at 60 kHz is limited, therefore it is not suitable to carry any voice or audio information. There is however, enough room to send code consisting of binary digits, or bits. These bits can only be one of two values: a zero or a one. Believe it or not, each bit is sent out at a relatively slow pace of one per second. In fact it takes up to a full minute just to complete time code containing the current date and time.

In case you are wondering how the NIST generates these bits being sent out from WWVB, it's by raising and lowering the power used to create the signal. NIST estimates its signal coverage based on the transmitted field strength of 100 microvolts per meter, which in theory should be large enough to guarantee decent reception in the continental United States. [19]

Radio clock receivers should be able tune into WWVB just about anywhere in North America with pretty good results. If you are in Hawaii you will need to tune into WWVH[20] which is broadcast on four different frequencies: 2.5 MHz, 5 mHz,10 MHz and 15 MHz. Europe, Asia or anywhere else in the world will not be able to receive the signal from these two NIST radio stations. Oddly enough, Alaska seems to be exempt from that last statement. (One theory is because there exist a distinct lack of radio frequency background noise found in sparsely populated regions.)

As with GPS, Radio receivers are not dissimilar when it comes to proper antenna placement/orientation and PC hardware installation. Also, remember that the accuracy of the final timestamp can be affected by the components that comprise the PC itself. Please the review earlier the portion of this document pertaining to GPS.

A decent radio receiver should be able to pick up signal with a lower sensitivity rating of about 50 microvolts per meter and produce a consistently accurate timestamp within 0.1 seconds. An example may be found at:

http://www.spectracomcorp.com/products/show_prod.php?id=24?source=overture&campaign=Antenna

Modem and Internet

These two devices are also available for synchronization with NIST strata 1 servers as previously mentioned in this paper under the “Windows and NIST application example”.

The following two descriptions are taken directly from the NIST website and can be found at: http://tf.nist.gov/timefreq/service/time-computer.html.

□ *By Internet*

This is the method of choice for most computer users. The NIST Internet Time Service (ITS) allows you to quickly synchronize the clock of any computer connected to the Internet. Simple client software allows you to synchronize your clock as often as necessary, and the service is completely free[21].

□ *By Telephone*

The NIST Automated Computer Time Service (ACTS) allows computers with analog modems to synchronize their clocks by telephone using simple client software. This service is intended for computers that are not connected to the Internet, or that are behind a firewall. It requires making a phone call of less than 1 minute each time you set your clock. The call is long distance outside of the Denver/Boulder, Colorado calling area[22].

Keep in mind that ACTS only works with analog modems that use ordinary telephone lines. Digital modems, such as Digital Subscriber Line (DSL) and cable modems, cannot connect to ACTS. If your computer has a digital modem, use the Internet Time Service to synchronize to NIST via your Internet connection[23].

End Times

So it appears that time synchronization is indeed moving into the future and will figure prominently as new technology is developed. Whether or not NTP will be able to keep pace will remain to be seen. One thing is for certain, it’s wide popularity has earned NTP a reputation worthy of the stars.

References

[1] Mills, David “A brief History of NTP Time: Confessions of an Internet Timekeeper.” May 25,2004 URL: http://www.ee.udel.edu/~mills/database/memos/hist.txt (September 21, 2004).

[2] Wray, J. “Request for Comments: 1509.” September 1993. URL: ftp://ftp.rfc-editor.org/in-notes/rfc1509.txt (September 23. 2004).

[3] Mills, David. “Network Time Protocol (Version 3) Specification, Implementation and Analysis.” March 1992. URL: ftp://ftp.rfc-editor.org/in-notes/rfc1305.txt (September 23. 2004).

[4] Mills, D. “RFC 2030 - Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI.” October 1996. URL: http://www.faqs.org/rfcs/rfc2030.html (September 23. 2004).

[5] Unknown. “Introduction to Public-Key Cryptography.” October 10 1998. URL: http://developer.netscape.com/docs/manuals/security/pkin (September 23. 2004).

[6] Unknown. “Introduction to Public-Key Cryptography.” October 10 1998. URL: http://developer.netscape.com/docs/manuals/security/pkin (September 23. 2004).

[7] Mills, David. “Public-Key Cryptography for the Network Time Protocol Version 1.” June 2000. URL: http://www.eecis.udel.edu/~mills/database/memos/draft-ietf-stime-ntpauth-00.txt (September 20, 2004).

[8] Mills, David. “NTP Version 4 Release Notes.” August 5, 2003. URL: http://www.eecis.udel.edu/~mills/ntp/html/release.html (September 20, 2004).

[9] Deeths, David. and Brunette, Glenn “Using NTP to Control and Synchronize Systems Clocks – Part I: Introduction to NTP” July, 2001. URL: http://www.sun.com/blueprints/0701/NTP.pdf (September 20, 2004).

[10] Mills, David “Public NTP Primary (stratum 1) Time Servers” July, 2001. URL: http://www.eecis.udel.edu/~mills/ntp/clock1b.html (September 20, 2004).

[11]U.S. Naval Observatory “What is Universal Time?” September 20, 2004. URL: http://aa.usno.navy.mil/faq/docs/UT.html (September 23, 2004).

[12] Mills, David "Public NTP Primary (stratum 1) Time Servers" August 24, 2004. URL: http://www.eecis.udel.edu/~mills/ntp/clock1b.html (September 20, 2004).

[13] Rash, James "Internet Technology on Spacecraft." September 2000 URL: http://ipinspace.gsfc.nasa.gov/documents/Space2000Paper-html (September 21, 2004).

[14] Microsoft "The Windows Time Service" April 27, 2001. URL: http://www.microsoft.com/windows2000/techinfo/howitworks/security/wintimeserv.asp (September 20, 2004).

[15] Ost, Linda "NIST Unveils Chip-Scale Atomic Clock" August 27, 2001. URL: http://groups.google.com/groups?q=atomic+clock+cesium+gps&hl=en&lr=&ie=UTF-8&c2coff=1&selm=413B5B50.2010005%40nova.astro.utoronto.ca&rnum=3 (September 21, 2004).

[16] Satellite Data Systems "http://www.sds-gps.com/glossaryc.html" URL: http://www.sds-gps.com/glossaryc.html (September 21, 2004).

[17] Bennett, Peter "The NMEA FAQ" June 12, 2003 URL: http://vancouver-webpages.com/peter/nmeafaq.txt (September 21, 2004).

[18] National Institute of Standards and Technology "NIST Radio Station WWVB" URL: http://tf.nist.gov/timefreq/stations/wwvb.htm (September 21, 2004).

[19] Beaglesoft "Frequently Asked Questions about GPS" September 13, 2004 URL: http://www.beaglesoft.com/gpsfaq.htm#howworksBeaglesoft (September 21, 2004).

[20] Wikipedia. "WWVH." June 10, 2004. URL: http://en.wikipedia.org/wiki/WWVH (September 20, 2004).

[21] National Institute of Standards and Technology "Set your Computer Clock to NIST TIME" URL: http://tf.nist.gov/timefreq/service/its.htm (September 21, 2004).

[22] National Institute of Standards and Technology "Set your Computer Clock to NIST TIME" URL: http://tf.nist.gov/timefreq/service/acts.htm (September 21, 2004).

[23] National Institute of Standards and Technology "Set your Computer Clock to NIST TIME" URL: http://tf.nist.gov/timefreq/service/its.htm (September 21, 2004).