White Paper

What Time is it?

The Importance of Time in the Network

Network Time Synchronization is becoming more important than ever in today's networks as companies have come to rely on their networks for all their communication needs. Since the 1990's we have seen an explosion of TCP/IP network usage and de-centralized network resources. Mainframes became PC's and having different networks for every application (TDM Voice, SNA Mainframe, video survelliance, building access, audio rendering, video rendering, IP Data Networks) is quickly being replaced with IP Networks running all of these applications as different services over the same network. The goal of having one network running many services, instead of a different disparate network per service, is a valid and more cost-effective approach to Enterprise communication needs.

IP Networks are now the primary communication medium for most Enterprise application needs including voice, wireless, data and video communications with the goal of being able to reduce the cost and number of different networks and systems, while at the same time, allowing applications and services increased access to data that can be quickly shared between each other.

Companies today increasingly seek network performance data and are no longer satisfied by basic IP connectivity alone. The questions about their data, applications, and services encompass not just 'if', and 'how much' of the data arrived, but exactly 'when' did it arrive? IP Networks now require Quality of Service (QoS), WAN Optimization, secured delivery and reciepts, traffic priorities, and more beyond basic IP connectivity services. All of these demands are beginning to drive a greater need for customers to be able to accurately measure network, service, and application performance, quality, and availibility.

This increased focus on network performance has placed new demands on the distribution of time within the network for many reasons. Log file correlation, connectivity test and measurement, event sequencing and forensics depend on time (Time of Day) system clock synchronization as a fundamental requirement.

Performance metrics such as One-Way (OW) and Round Trip (RTT) network latencies, VoIP Mean Opinion Scores (MoS), packet loss, packet reordering, and packet interarrival delay (jitter), data compression, and queue delay are watched every day with great interest by network operations groups. Most of these metrics and data have a timestamp, or time of day information associated with them, especially single-path one-way delay. It is very common in networks to have a different traffic path on both parts of a round-trip, which makes using round-trip metrics alone not as useful as one-way metrics in these scenarios.

One-Way and Round-Trip Latecy

Many TCP/IP operations and data are largely One-Way traffic by nature such as video, Multicast, and even traditional data protocols such as FTP, Web/HTTP, Network File and Print Services, their primarily function is to deliver most of their packet traffic from Server to Client. Traditional voice (VoIP) traffic is largely bi-directional in nature and not generally classified as One-Way data though each party to a call has their own transmission and reception traffic that are largely one-way in nature. Audio broadcast or multicast traffic is very much own-way by design and financial markets use a tremendous amount of one-way data transfers in the form of equity markets realtime data publishing.

The basic requirement to have the network systems clocks synchronized to a common reference clock is complicated by the very fact that IP network signals and transmissions were not designed as synchronized or deterministic for this purpose, which helped reduce their costs and complexity. SONET, ATM, Token-Ring, and FDDI are common examples of networks that are all fairly well synchronized but are less common today. As we wish to distribute time over the very same IP networks that we are trying to measure, the nature of these networks introduces errors in the distribution of time as its data transverses over the network. As the The Time and Frequency Division of the National Institute of Standards and Technology mentions at their “NIST Research on Digital Time Services” website stated, “The accuracy of the digital time messages is limited primarily by the jitter in the delay through the transmission network,..” [6].

What Can be Done?

This paper does not intend to be a technical deep dive into the intricate layers of Time and Frequency Synchronization, but more simply intends to primarily give a glimpse into why this is such a complex subject, why the Network Time Protocol (NTP), which is so pervasive and commonly used everywhere in IP networks and systems today, cannot solve this issue alone, and what else you may be able to do to help address the problem.

If you could query every device in the network instantaneously and ask what the time of day each device thinks it is, they will generally answer back much like any computer and will report some differences between all the devices ranging from many seconds to a minute or more, even on a network operating well.

Why is this so difficult, you ask? Due to time itself: the time it takes to get time and the delay and jitter involved in the transfer of time over the IP Network. It is an elusive subject and once we attempt to dig deeper and uncover the underlying concepts of Time Accuracy and Clock Synchronization, what originally seemed like a casual little issue of a few seconds, now turns into a Grand Canyon size issue of 100’s of milliseconds and 1000’s of microseconds.

NTP is non-deterministic and uncontrolled under general network operating conditions and was designed to be more common and easy to implement across many types of computing devices. Other alternatives, like directly connected GPS which require an antennae and line of sight visibility to the sky, are not always practical or cost-effective. In telelcom Central Offices (CO) and Cable Headends there are other synchronization technologies like SONET, SDH and Synchronous Ethernet that can deliver frequency, but not time. There is also the DOCSIS Timing Interace (DTI) and Universal Timing Interface (UTI) which can deliver very preceise nanosecond time accuracy, but are limited to use within a CO or Headend. These technologies can co-exist and compliment NTP’s performance. In some manufacturing environments there is an “IEEE-1588 Standard for Synchronizing Clocks” [7] time synchronization technology that can be accurate to sub-microseconds and works well in controlled industrial

LAN environments. Yet IEEE-1588 is not well suited to WAN environments, requires direct hardware technology, and does not work with packet encapsulation or encryption built on top, and it also lacks typical carrier-class service features like Authentication, and control paremeters similar to those that exist in NTP today.

An IEEE-1588 deployment naturally requires a new architecture of Servers and Cleints and does not address the install base of NTP applications. There are efforts in the IETF to improve NTP with similar techniques, without changing the protocol itself. Today's NTP protocol is the best method to distribute time across the network, however it's deployment acrhitecture needs improvement to meet the demands of new applications.

It Takes Time to Update Time

What is the heart of the problem for Time Synchronization? 1.) Clock Accuracy, or the ability of a hardware clock to prevent deviation and drift once set to a particular time, and 2.) Synchronization or the ability of a local clock to get time from a reference source such as US Naval Observatory's Master Clock in Washington DC.

Clock physical accuracy and stability boils down to cost. Clocks made with extremely stable materials (Rubidium, Cesium, etc) have the ability to hold better time once set, versus clocks made from less expensive and more commonly available materials. Clock synchronization is an even less deterministic subject as the problem domain is not contained to something as simple as materials. Synchronization requires protocols, algorithms, and estimations to negotiate with another clock or a time source in another location. The local clock, be it cesium or silicon, requesting a time update must make a great deal of assumptions between when it sends out that update request and when it receives the response back. This may sound odd, but the local clock must estimate how long the time request took (elapsed time) broken into two fundamental parts: A.) how long did it take the request to get to the source reference clock, and B.) how long did it take the time data to get back to the requesting local clock. Lastly, the process, the software or hardware that requested the time update is responsible for updating its own local clock and it must assume there is zero delay in the update to its local clock; that is the time it takes to update itself.

If this sounds humorous to you then you can be assured you are not alone. Such conversations about time accuracy by anyone not in the time accuracy business can get a little interesting as we talk about time knowing about time, yet even small errors in any of the components or protocols mentioned will add up very quickly to errors in your computers clock time. A few milliseconds here and there will quickly add up to a second. One hundred milliseconds here and there will quickly add up to 10 seconds.

One of the most complex issues regards the fact that your local clock and time update service can only estimate how far off it may be from the reference clock time source, which makes compensation algorithms inherently only nominally accurate. Without the addition of a directly connected device hardware clock pulse like direct GPS plus PPS, or UTI (Universal Timing Interface) used in telecom Central Offices (COs), these estimations can only be based on statistical assumptions which due to the errors in the their data source and how they get their data, will always be incorrect by some order of magnitude.

Network Time Protocol (NTP)

As IP and NGN Networks begin to support more and more business services and applications like voice, video and increasing amounts of data, network operators continue to look more and more closely at time distribution and accuracy across the network. In the past it was assumed that the Network Time Protocol (NTP) was generally sufficient for such purposes, yet as operators ask harder questions related to time such as 'what is my one-way latency between these two points' in the network, they are finding NTP is, and never really was as deterministic as we all would like it to be. So, how good is NTP and how do we improve it? And, where are its limits or reliability and accuracy where another directly connected hardware interface similar to GPS, UTI, DTI might be needed?

NTP History and Background

Network Time Protocol (NTP) is an Internet protocol for synchronizing system clocks among a set of distributed time servers and clients across the network to standard time. NTP was initially invented to provide time to computer hosts in both private networks and the public Internet. Gradually NTP's use grew and today it is pervasive and is the universally accepted method to synchronize system clocks of computers, servers and data communication equipment.

Historically, NTP has been deployed in the telecom IT departments or data centers for post processing functions, typically known as back office to support operational activities such as billing and event log generation. Today, new applications that require NTP have emerged both from a functional perspective as well as service perspective. Some new applications, like Service Level Agreements (SLA's), require a more accurate and assured NTP than can be provided by the systems in place. In general, there are three types of NTP Server devices: a primary server, secondary server, and a client. A primary server is synchronized directly to a reference clock, such as a GPS receiver. A secondary server has one or more upstream servers (primary or secondary) and one or more downstream secondary servers or clients. A client is synchronized to one or more upstream servers, but does not provide synchronization to dependent clients.

Stratum Levels in NTP

The flexibility of NTP's architecture has made it simple to deploy and implement. It however has also led to a high degree of variability in performance. Since most applications historically did not require very much accuracy, the primary goal of NTP when it was designed was to ensure that NTP was always available from one or more servers. In NTP the one primary metrics that is used to describe how NTP Servers and Clients are connected is the term Stratum. Stratum simply refers to a logical heirachial arrangement of multiple NTP servers. In essence, the more Stratum levels a client or server is away from a Stratum 1 Server, the less accurate it will be, though again knowning exactly how inaccurate any NTP Server or Client may be, except for one directly physically connected to a reference clock (Stratum 1), is still just as complicated and unknown. It is generally accepted that when it is possible, having an NTP Server or Client getting it's reference time from a Stratum 1 NTP Server vs. a Stratum 2 or lower NTP Server will give better accuracy.

Figure 1. NTP Hierarchy (Strata Model)

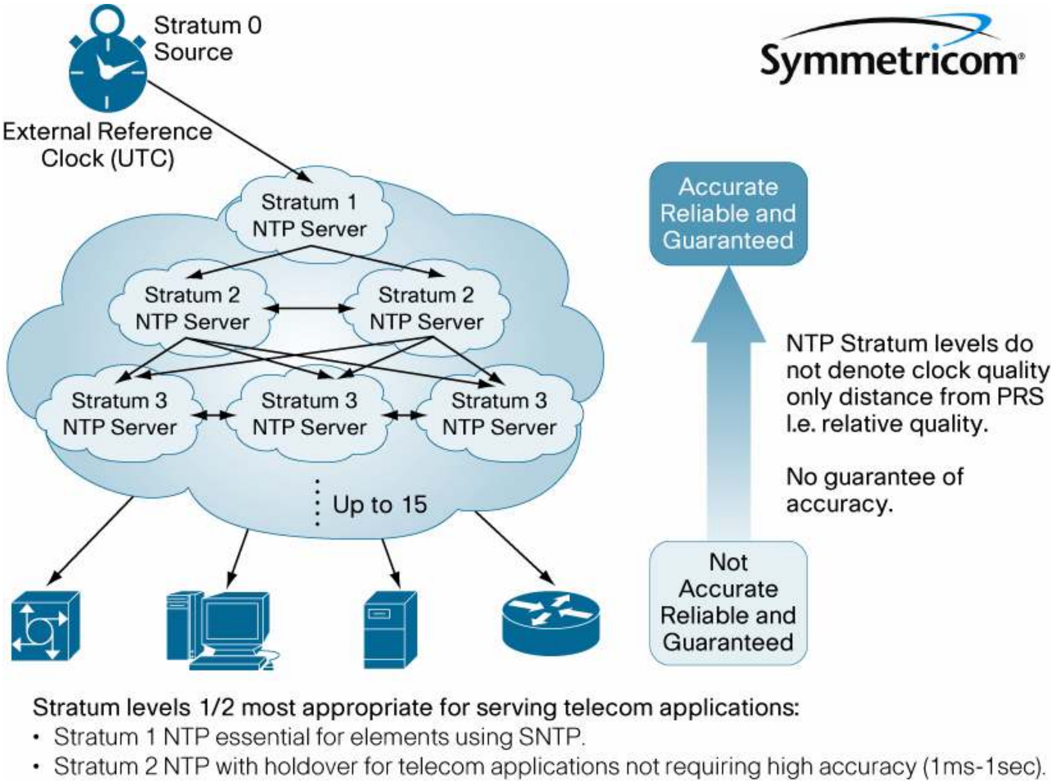

Stratum levels 1/2 most appropriate for serving telecom applications:

- Stratum 1 NTP essential for elements using SNTP.
- Stratum 2 NTP with holdover for telecom applications not requiring high accuracy (1ms-1sec).

A Stratum 1 NTP Server is a primary server connected to what is known as a reference clock, which is typically a GPS clock. If a secondary NTP Server is then connected to the Stratum 1 Server it now becomes a Stratum 2 server. Another secondary NTP Server connected to that server would then be a Stratum 3 Server. The Stratum metric is useful in understanding how NTP is distributed in the network, but does not guarantee accuracy. Even two Stratum 1 NTP servers may have very different performance, since they may be designed and implemented differently. One of the uses of the Stratum information is when a Server fails. For example, a client may be connected to a Stratum 4 NTP Server. If that server fails it may use the Stratum information to find another Stratum 4 NTP Server or better in the network. However, to ensure performance of NTP during a failure, it is important to ensure that NTP servers are deployed with similar accuracy. One way to ensure accuracy is maintained during failure is to flatten the NTP deployment architecutre by placing NTP Servers near the edge of the network. Moreover, avoiding failure itself can be done by protecting the common equipment in the NTP Server to achive high availability (reliability).

With NTP widely deployed and used by many applications, the next step is to address applications that need accuracy, in addition to high availability. The first step in supporting accurate NTP is to establish a baseline accuracy that is common across several points in the network. Once a baseline accuracy is established and known by the applications that depend on it, then the accuracy can be improved. For example, a profile of NTP performance can be established to ensure that if an NTP Server is lost, another could be used with known performance bounds.

Using a carrier-class NTP Server that has a guaranteed performance is the first step in ensuring accurate NTP. The second step is to ensure that the Server is deployed near the edge of the network. The network will impose variability on the NTP packets and thus reduce the performance. The closer the Server and Clients are, the easier it is to ensure a baseline accuracy.

The first step towards establishing a baseline accuracy for NTP is choosing a carrier-class NTP server, rather than using a simple and non-deterministic Linux or Unix software based NTP server system. Such NTP servers differentiate themselves from the average server in several ways. A carrier-class NTP Server ensures that the NTP timestamps leaving the server have a consistent and common accuracy across all server-client relationships, under various loading factors and over long time periods. In addition to having consistent high accuracy, a carrier-class NTP server has several redundancy mechanisms within the server itself and also through peering with other NTP servers. Lastly, a carrier-class server offers extensive manageability and performance metrics for the servers operation and each client that uses it.

The third aspect to increase NTP accuracy is to baseline and improve the performance of the NTP Client. Since NTP is typically implemented in software, there are several factors that impact performance including processor speed, queuing, etc. In the future, some NTP clients may be implemented within systems hardware instead of process-level software, which can help improve accuracy considerably.

NTP in Cisco IOS Software

Cisco Internet Operating System (IOS) Software currently runs the standard implementation of Network Time Protocol (NTP) v3.0. NTP is used as the general method for time synchronization as described above. While NTP has generally been assumed to be accurate enough for most event logging and event forensics across network devices, even those applications that consume time information are being pressed for more accurate synchronization. Other applications like Cisco IOS IP SLA’s are being asked to provide greater data accuracy, which in turn relies upon and requires greater NTP time synchronization accuracy.

Cisco IOS IP SLA's

Cisco IOS IP SLA’s is a IOS embedded network performance and health management feature that can be used to perform point to point TCP/IP, VoIP, MPLS, and Metro-Ethernet tests. These test operations are designed to measure and notify in realtime on network latency, jitter, packet and frame loss, elapsed times, and connectivity or path availiblity. IP SLA’s also support basic network application service tests like HTTP, DNS, DHCP, TCP Connect, and FTP.

Many of these IP SLAs test operations measure both One-Way and Round-Trip elapsed times returning data about network latency or delay. One key latency metric who’s accuracy is directly linked to NTP accuracy is our UDP Jitter One-Way Delay test operation. This IP SLAs operation can test and report the OW delay for both the outbound and return paths in each single operation between two points in the network. Here, the data returned by IP SLAs will be directly linked to the engineering design, deployment, and accuracy of NTP in your networks.

Cisco and Symmetricom

For More Information

If you have any questions about IOS IP SLA’s and wish to have some NTP best-practices design and analysis help, we suggest you contact Cisco and Symmetricom directly. Symmetricom are experts in the fields of Time and Frequency Synchronization and are the supplier of choice for our Cisco IP SLAs Engineering Teams Network Time Protocol (NTP) Servers.

References

[1] UTC DISSEMINATION TO THE REAL-TIME USER: THE ROLE OF USNO, Mihran Miranian U.S. Naval Observatory, Washington, D.C. 20392, http://tycho.usno.navy.mil/ptti/1995/Vol%2027_05.pdf

[2] Network Time Protocol, Wikipedia.org, http://en.wikipedia.org/wiki/Network_Time_Protocol

[3] David L. Mills, PHD, Professor, University of Delaware, http://www.ee.udel.edu/~mills/

[4] NTP Home Page, http://www.ntp.org/

[5] NIST Internet Time Service, http://tf.nist.gov/service/its.htm

[6] NIST, "NIST Research on Digital Time Services", http://tf.nist.gov/timefreq/time/digitaltime.htm

[7] NIST, "IEEE 1588-2002 Standard for Synchronizing Clocks", http://ieee1588.nist.gov/

Related Technologies and Links

[8] Universl Timing Interface (UTI), http://www.itu.int/md/T05-SG09-051017-D-0038/en

[9] Docsis Timing Interface (DTI), http://www.cablemodem.com/downloads/specs/CM-SP-DTI-I04-061222.pdf

[10] United States Naval Observatory, http://www.usno.navy.mil/

[11] Network Time Protocol [NTP], http://www.ntp.org/index.html

[12] Internet Engineering Task Force (IETF), http://www3.ietf.org/

[13] IETF NTP Working Group, http://www.ietf.org/html.charters/ntp-charter.html

For More Information

For more information about the Cisco IP SLA's, visit http://www.cisco.com/go/ipsla or contact your local account representative.

For more information about Symmetricom, visit http://www.symmetricom.com/.

For more information about the Cisco End-of-Life Policy, go to http://www.cisco.com/en/US/products/prod_end_of_life.html.

To subscribe to receive end-of-life/end-of-sale information, go to http://www.cisco.com/cgi-bin/Support/FieldNoticeTool/field-notice.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-425832-01 05/13