

Network Working Group
Request for Comments: 1769
Obsoletes: 1361
Category: Informational

D. Mills
University of Delaware
March 1995

Simple Network Time Protocol (SNTP)

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This memorandum describes the Simple Network Time Protocol (SNTP), which is an adaptation of the Network Time Protocol (NTP) used to synchronize computer clocks in the Internet. SNTP can be used when the ultimate performance of the full NTP implementation described in RFC-1305 is not needed or justified. It can operate in both unicast modes (point to point) and broadcast modes (point to multipoint). It can also operate in IP multicast mode where this service is available. SNTP involves no change to the current or previous NTP specification versions or known implementations, but rather a clarification of certain design features of NTP which allow operation in a simple, stateless remote-procedure call (RPC) mode with accuracy and reliability expectations similar to the UDP/TIME protocol described in RFC-868.

This memorandum obsoletes RFC-1361 of the same title. Its purpose is to explain the protocol model for operation in broadcast mode, to provide additional clarification in some places and to correct a few typographical errors. A working knowledge of the NTP Version 3 specification RFC-1305 is not required for an implementation of SNTP. Distribution of this memorandum is unlimited.

1. Introduction

The Network Time Protocol (NTP) specified in RFC-1305 [MIL92] is used to synchronize computer clocks in the global Internet. It provides comprehensive mechanisms to access national time and frequency dissemination services, organize the time-synchronization subnet and adjust the local clock in each participating subnet peer. In most places of the Internet of today, NTP provides accuracies of 1-50 ms, depending on the characteristics of the synchronization source and network paths.

RFC-1305 specifies the NTP protocol machine in terms of events, states, transition functions and actions and, in addition, optional algorithms to improve the timekeeping quality and mitigate among several, possibly faulty, synchronization sources. To achieve accuracies in the low milliseconds over paths spanning major portions of the Internet of today, these intricate algorithms, or their functional equivalents, are necessary. However, in many cases accuracies of this order are not required and something less, perhaps in the order of large fractions of the second, is sufficient. In such cases simpler protocols such as the Time Protocol [POS83], have been used for this purpose. These protocols usually involve an RPC exchange where the client requests the time of day and the server returns it in seconds past some known reference epoch.

NTP is designed for use by clients and servers with a wide range of capabilities and over a wide range of network delays and jitter characteristics. Most users of the Internet NTP synchronization subnet of today use a software package including the full suite of NTP options and algorithms, which are relatively complex, real-time applications. While the software has been ported to a wide variety of hardware platforms ranging from supercomputers to personal computers, its sheer size and complexity is not appropriate for many applications. Accordingly, it is useful to explore alternative access strategies using far simpler software appropriate for less stringent accuracy expectations.

This memorandum describes the Simple Network Time Protocol (SNTP), which is a simplified access strategy for servers and clients using NTP as now specified and deployed in the Internet. There are no changes to the protocol or implementations now running or likely to be implemented in the near future. The access paradigm is identical to the UDP/TIME Protocol and, in fact, it should be easily possible to adapt a UDP/TIME client implementation, say for a personal computer, to operate using SNTP. Moreover, SNTP is also designed to operate in a dedicated server configuration including an integrated radio clock. With careful design and control of the various latencies in the system, which is practical in a dedicated design, it is possible to deliver time accurate to the order of microseconds.

It is strongly recommended that SNTP be used only at the extremities of the synchronization subnet. SNTP clients should operate only at the leaves (highest stratum) of the subnet and in configurations where no NTP or SNTP client is dependent on another SNTP client for synchronization. SNTP servers should operate only at the root (stratum 1) of the subnet and then only in configurations where no other source of synchronization other than a reliable radio clock is available. The full degree of reliability ordinarily expected of primary servers is possible only using the redundant sources, diverse

subnet paths and crafted algorithms of a full NTP implementation. This extends to the primary source of synchronization itself in the form of multiple radio clocks and backup paths to other primary servers should the radio clock fail or deliver incorrect time. Therefore, the use of SNTP rather than NTP in primary servers should be carefully considered.

2. Operating Modes and Addressing

Like NTP, SNTP can operate in either unicast (point to point) or broadcast (point to multipoint) modes. A unicast client sends a request to a server and expects a reply from which it can determine the time and, optionally, the roundtrip delay and local clock offset relative to the server. A broadcast server periodically sends a message to a designated IP broadcast address or IP multicast group address and ordinarily expects no requests from clients, while a broadcast client listens on this address and ordinarily sends no requests to servers. Some broadcast servers may elect to respond to client requests as well as send unsolicited broadcast messages, while some broadcast clients may elect to send requests only in order to determine the network propagation delay between the server and client.

In unicast mode the client and server IP addresses are assigned following the usual conventions. In broadcast mode the server uses a designated IP broadcast address or IP multicast group address, together with a designated media-access broadcast address, and the client listens on these addresses. For this purpose, an IP broadcast address has scope limited to a single IP subnet, since routers do not propagate IP broadcast datagrams. In the case of Ethernets, for example, the Ethernet media-access broadcast address (all ones) is used with an IP address consisting of the IP subnet number in the net field and all ones in the host field.

On the other hand, an IP multicast group address has scope extending to potentially the entire Internet. The actual scope, group membership and routing are determined by the Internet Group Management Protocol (IGMP) [DEE89] and various routing protocols, which are beyond the scope of this document. In the case of Ethernets, for example, the Ethernet media-access broadcast address (all ones) is used with the assigned IP multicast group address of 224.0.1.1. Other than the IP addressing conventions and IGMP, there is no difference in server operations with either the IP broadcast address or IP multicast group address.

Broadcast clients listen on the designated media-access broadcast address, such as all ones in the case of Ethernets. In the case of IP broadcast addresses, no further provisions are necessary. In the case

of IP multicast group addresses, the host may need to implement IGMP in order that the local router intercepts messages to the 224.0.1.1 multicast group. These considerations are beyond the scope of this document.

In the case of SNTP as specified herein, there is a very real vulnerability that SNTP multicast clients can be disrupted by misbehaving or hostile SNTP or NTP multicast servers elsewhere in the Internet, since at present all such servers use the same IP multicast group address 224.0.1.1. Where necessary, access control based on the server source address can be used to select only those servers known to and trusted by the client. Alternatively, by convention and informal agreement, all NTP multicast servers now include an MD5-encrypted message digest in every message, so that clients can determine if the message is authentic and not modified in transit. It is in principle possible that SNTP clients could implement the necessary encryption and key-distribution schemes, but this is considered not likely in the simple systems for which SNTP is intended.


While not integral to the SNTP specification, it is intended that IP broadcast addresses will be used primarily in IP subnets and LAN segments including a fully functional NTP server with a number of SNTP clients in the same subnet, while IP multicast group addresses will be used only in special cases engineered for the purpose. In particular, IP multicast group addresses should be used in SNTP servers only if the server implements the NTP authentication scheme described in RFC-1305, including support for the MD5 message-digest algorithm.

3. NTP Timestamp Format

SNTP uses the standard NTP timestamp format described in RFC-1305 and previous versions of that document. In conformance with standard Internet practice, NTP data are specified as integer or fixed-point quantities, with bits numbered in big-endian fashion from 0 starting at the left, or high-order, position. Unless specified otherwise, all quantities are unsigned and may occupy the full field width with an implied 0 preceding bit 0.

Since NTP timestamps are cherished data and, in fact, represent the main product of the protocol, a special timestamp format has been established. NTP timestamps are represented as a 64-bit unsigned fixed-point number, in seconds relative to 0h on 1 January 1900. The integer part is in the first 32 bits and the fraction part in the last 32 bits. In the fraction part, the non-significant low-order bits should be set to 0. This format allows convenient multiple-precision arithmetic and conversion to UDP/TIME representation

(seconds), but does complicate the conversion to ICMP Timestamp message representation (milliseconds). The precision of this representation is about 200 picoseconds, which should be adequate for even the most exotic requirements.

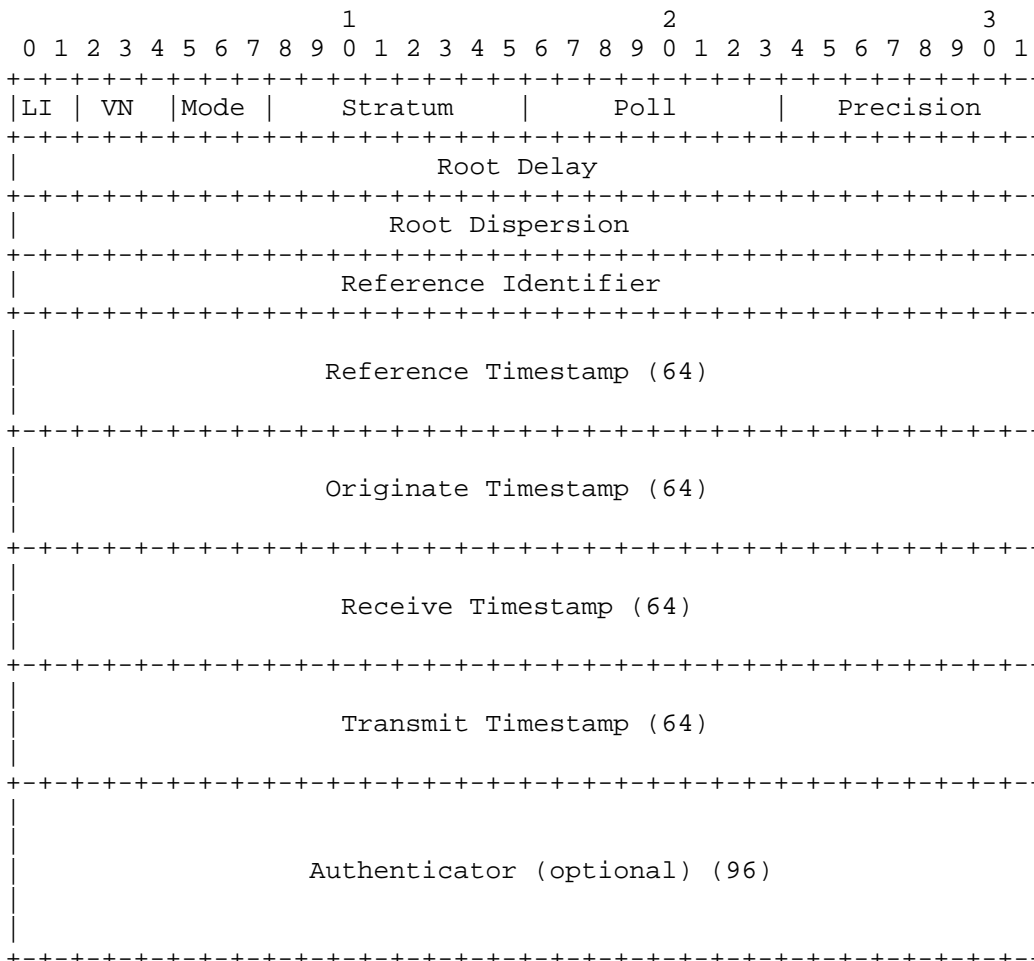


Note that, since some time in 1968 the most significant bit (bit 0 of the integer part) has been set and that the 64-bit field will overflow some time in 2036. Should NTP or SNTP be in use in 2036, some external means will be necessary to qualify time relative to 1900 and time relative to 2036 (and other multiples of 136 years). Timestamped data requiring such qualification will be so precious that appropriate means should be readily available. There will exist a 200-picosecond interval, henceforth ignored, every 136 years when the 64-bit field will be 0, which by convention is interpreted as an invalid or unavailable timestamp.

4. NTP Message Format

Both NTP and SNTP are clients of the User Datagram Protocol (UDP) [POS80], which itself is a client of the Internet Protocol (IP) [DAR81]. The structure of the IP and UDP headers is described in the cited specification documents and will not be described further here. The UDP port number assigned to NTP is 123, which should be used in both the Source Port and Destination Port fields in the UDP header. The remaining UDP header fields should be set as described in the specification.

Following is a description of the SNTP message format, which follows the IP and UDP headers. The SNTP message format is identical to the NTP format described in RFC-1305, with the exception that some of the data fields are "canned," that is, initialized to pre-specified values. The format of the NTP message is shown below.



As described in the next section, in SNTP most of these fields are initialized with pre-specified data. For completeness, the function of each field is briefly summarized below.

Leap Indicator (LI): This is a two-bit code warning of an impending leap second to be inserted/deleted in the last minute of the current day, with bit 0 and bit 1, respectively, coded as follows:

LI	Value	Meaning
00	0	no warning
01	1	last minute has 61 seconds
10	2	last minute has 59 seconds)
11	3	alarm condition (clock not synchronized)

Version Number (VN): This is a three-bit integer indicating the NTP version number, currently 3.

Mode: This is a three-bit integer indicating the mode, with values defined as follows:

Mode	Meaning
0	reserved
1	symmetric active
2	symmetric passive
3	client
4	server
5	broadcast
6	reserved for NTP control message
7	reserved for private use

In unicast mode the client sets this field to 3 (client) in the request and the server sets it to 4 (server) in the reply. In broadcast mode the server sets this field to 5 (broadcast).

Stratum: This is a eight-bit unsigned integer indicating the stratum level of the local clock, with values defined as follows:

Stratum	Meaning
0	unspecified or unavailable
1	primary reference (e.g., radio clock)
2-15	secondary reference (via NTP or SNTP)
16-255	reserved

Poll Interval: This is an eight-bit signed integer indicating the maximum interval between successive messages, in seconds to the nearest power of two. The values that can appear in this field presently range from 4 (16 s) to 14 (16284 s); however, most applications use only the sub-range 6 (64 s) to 10 (1024 s).

Precision: This is an eight-bit signed integer indicating the precision of the local clock, in seconds to the nearest power of two. The values that normally appear in this field range from -6 for mains-frequency clocks to -20 for microsecond clocks found in some workstations.

Root Delay: This is a 32-bit signed fixed-point number indicating the total roundtrip delay to the primary reference source, in seconds with fraction point between bits 15 and 16. Note that this variable can take on both positive and negative values, depending on the relative time and frequency offsets. The values that normally appear in this field range from negative values of a few milliseconds to positive values of several hundred milliseconds.

Root Dispersion: This is a 32-bit unsigned fixed-point number indicating the nominal error relative to the primary reference source, in seconds with fraction point between bits 15 and 16. The values that normally appear in this field range from 0 to several hundred milliseconds.

Reference Clock Identifier: This is a 32-bit code identifying the particular reference source. In the case of stratum 0 (unspecified) or stratum 1 (primary reference), this is a four-octet, left-justified, 0-padded ASCII string. While not enumerated as part of the NTP specification, the following are representative ASCII identifiers:

Stratum	Code	Meaning
1	pps	precision calibrated source, such as ATOM (atomic clock), PPS (precision pulse-per-second source), etc.
1	service	generic time service other than NTP, such as ACTS (Automated Computer Time Service), TIME (UDP/Time Protocol), TSP (Unix Time Service Protocol), DTSS (Digital Time Synchronization Service), etc.
1	radio	Generic radio service, with callsigns such as CHU, DCF77, MSF, TDF, WWV, WWVB, WWVH, etc.
1	nav	radionavigation system, such as OMEG (OMEGA), LORC (LORAN-C), etc.
1	satellite	generic satellite service, such as GOES (Geostationary Orbit Environment Satellite, GPS (Global Positioning Service), etc.
2	address	secondary reference (four-octet Internet address of the NTP server)

Reference Timestamp: This is the time at which the local clock was last set or corrected, in 64-bit timestamp format.

Originate Timestamp: This is the time at which the request departed the client for the server, in 64-bit timestamp format.

Receive Timestamp: This is the time at which the request arrived at the server, in 64-bit timestamp format.

Transmit Timestamp: This is the time at which the reply departed the server for the client, in 64-bit timestamp format.

Authenticator (optional): When the NTP authentication mechanism is implemented, this contains the authenticator information defined in Appendix C of RFC-1305. In SNTP this field is ignored for incoming messages and is not generated for outgoing messages.

5. SNTP Client Operations

The model for a SNTP client operating with either a NTP or SNTP server is a RPC client with no persistent state. In unicast mode, the client sends a client request (mode 3) to the server and expects a server reply (mode 4). In broadcast mode, the client sends no request and waits for a broadcast message (mode 5) from one or more servers, depending on configuration. Unicast client and broadcast server messages are normally sent at periods from 64 s to 1024 s, depending on the client and server configurations.

A unicast client initializes the SNTP message header, sends the message to the server and strips the time of day from the reply. For this purpose all of the message-header fields shown above are set to 0, except the first octet. In this octet the LI field is set to 0 (no warning) and the Mode field is set to 3 (client). The VN field must agree with the software version of the NTP or SNTP server; however, NTP Version 3 (RFC-1305) servers will also accept Version 2 (RFC-1119) and Version 1 (RFC-1059) messages, while NTP Version 2 servers will also accept NTP Version 1 messages. Version 0 (RFC-959) messages are no longer supported. Since there are NTP servers of all three versions interoperating in the Internet of today, it is recommended that the VN field be set to 1.

In both unicast and broadcast modes, the unicast server reply or broadcast message includes all the fields described above; however, in SNTP only the Transmit Timestamp has explicit meaning and then only if nonzero. The integer part of this field contains the server time of day in the same format as the UDP/TIME Protocol [POS83]. While the fraction part of this field will usually be valid, the accuracy achieved with SNTP may justify its use only to a significant

fraction of a second. If the Transmit Timestamp field is 0, the message should be disregarded.

In broadcast mode, a client has no additional information to calculate the propagation delay between the server and client, as the Transmit Timestamp and Receive Timestamp fields have no meaning in this mode. Even in unicast mode, most clients will probably elect to ignore the Originate Timestamp and Receive Timestamp fields anyway. However, in unicast mode a simple calculation can be used to provide the roundtrip delay d and local clock offset t relative to the server, generally to within a few tens of milliseconds. To do this, the client sets the Originate Timestamp in the request to the time of day according to its local clock converted to NTP timestamp format. When the reply is received, the client determines a Destination Timestamp as the time of arrival according to its local clock converted to NTP timestamp format. The following table summarizes the four timestamps.

Timestamp Name	ID	When Generated
Originate Timestamp	T1	time request sent by client
Receive Timestamp	T2	time request received at server
Transmit Timestamp	T3	time reply sent by server
Destination Timestamp	T4	time reply received at client

The roundtrip delay d and local clock offset t are defined as

$$d = (T4 - T1) - (T2 - T3)$$

$$t = ((T2 - T1) + (T3 - T4)) / 2.$$

The following table is a summary of the SNTP client operations. There are two recommended error checks shown in the table. In all NTP versions, if the LI field is 3, or the Stratum field is not in the range 1-15, or the Transmit Timestamp is 0, the server has never synchronized or not synchronized to a valid timing source within the last 24 hours. At the client discretion, the values of the remaining fields can be checked as well. Whether to believe the transmit timestamp or not in case one or more of these fields appears invalid is at the discretion of the implementation.

Field Name	Request	Reply
LI	0	leap indicator; if 3 (unsynchronized), disregard message
VN	1 (see text)	ignore
Mode	3 (client)	ignore
Stratum	0	ignore
Poll	0	ignore
Precision	0	ignore
Root Delay	0	ignore
Root Dispersion	0	ignore
Reference Identifier	0	ignore
Reference Timestamp	0	ignore
Originate Timestamp	0 (see text)	ignore (see text)
Receive Timestamp	0	ignore (see text)
Transmit Timestamp	0	time of day; if 0 (unsynchronized), disregard message
Authenticator	(not used)	ignore

6. SNTP Server Operations

The model for a SNTP server operating with either a NTP or SNTP client is an RPC server with no persistent state. Since a SNTP server ordinarily does not implement the full set of NTP algorithms intended to support redundant peers and diverse network paths, it is recommended that a SNTP server be operated only in conjunction with a source of external synchronization, such as a reliable radio clock. In this case the server always operates at stratum 1.

A server can operate in unicast mode, broadcast mode or both at the same time. In unicast mode the server receives a request message, modifies certain fields in the NTP or SNTP header, and returns the message to the sender, possibly using the same message buffer as the request. The server may or may not respond if not synchronized to a correctly operating radio clock, but the preferred option is to respond, since this allows reachability to be determined regardless of synchronization state. In unicast mode, the VN and Poll fields of the request are copied intact to the reply. If the Mode field of the request is 3 (client), it is set to 4 (server) in the reply; otherwise, this field is set to 2 (symmetric passive) in order to conform to the NTP specification.

In broadcast mode, the server sends messages only if synchronized to a correctly operating reference clock. In this mode, the VN field is set to 3 (for the current SNTP version), and the Mode field to 5 (broadcast). The Poll field is set to the server poll interval, in

seconds to the nearest power of two. It is highly desirable that, if a server supports broadcast mode, it also supports unicast mode. This is necessary so a potential broadcast client can calculate the propagation delay using client/server messages prior to regular operation using only broadcast messages.

The remaining fields are set in the same way in both unicast and broadcast modes. Assuming the server is synchronized to a radio clock or other primary reference source and operating correctly, the Stratum field is set to 1 (primary server) and the LI field is set to 0; if not, the Stratum field is set to 0 and the LI field is set to 3. The Precision field is set to reflect the maximum reading error of the local clock. For all practical cases it is computed as the negative of the number of significant bits to the right of the decimal point in the NTP timestamp format. The Root Delay and Root Dispersion fields are set to 0 for a primary server; optionally, the Root Dispersion field can be set to a value corresponding to the maximum expected error of the radio clock itself. The Reference Identifier is set to designate the primary reference source, as indicated in the table above.

The timestamp fields are set as follows. If the server is unsynchronized or first coming up, all timestamp fields are set to zero. If synchronized, the Reference Timestamp is set to the time the last update was received from the radio clock or, if unavailable, to the time of day when the message is sent. The Receive Timestamp and Transmit Timestamp fields are set to the time of day when the message is sent. In unicast mode, the Originate Timestamp field is copied unchanged from the Transmit Timestamp field of the request. It is important that this field be copied intact, as a NTP client uses it to check for replays. In broadcast mode, this field is set to the time of day when the message is sent. The following table summarizes these actions.

Field Name	Request	Reply
LI	ignore	0 (normal), 3 (unsynchronized)
VN	1, 2 or 3	3 or copied from request
Mode	3 (see text)	2, 4 or 5 (see text)
Stratum	ignore	1 server stratum
Poll	ignore	copied from request
Precision	ignore	server precision
Root Delay	ignore	0
Root Dispersion	ignore	0 (see text)
Reference Identifier	ignore	source identifier
Reference Timestamp	ignore	0 or time of day
Originate Timestamp	ignore	0 or time of day or copied from Transmit Timestamp of request
Receive Timestamp	ignore	0 or time of day
Transmit Timestamp	(see text)	0 or time of day
Authenticator	ignore	(not used)

There is some latitude on the part of most clients to forgive invalid timestamps, such as might occur when first coming up or during periods when the primary reference source is inoperative. The most important indicator of an unhealthy server is the LI field, in which a value of 3 indicates an unsynchronized condition. When this value is displayed, clients should discard the server message, regardless of the contents of other fields.

7. References

[DAR81] Postel, J., "Internet Protocol - DARPA Internet Program Protocol Specification", STD 5, RFC 791, DARPA, September 1981.

[DEE89] Deering, S., "Host Extensions for IP Multicasting. STD 5, RFC 1112, Stanford University, August 1989.

[MIL92] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation and Analysis. RFC 1305, University of Delaware, March 1992.

[POS80] Postel, J., "User Datagram Protocol", STD 6, RFC 768, USC/Information Sciences Institute, August 1980.

[POS83] Postel, J., and K. Harrenstien, "Time Protocol", STD 26, RFC 868, USC/Information Sciences Institute, SRI, May 1983.

Security Considerations

Security issues are not discussed in this memo.

Author's Address

David L. Mills
Electrical Engineering Department
University of Delaware
Newark, DE 19716

Phone: (302) 831-8247
EMail: mills@udel.edu