Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: September 15, 2016

H. Stenn
Network Time Foundation

March 14, 2016

Network Time Protocol Leap Smear REFID
draft-stenn-ntp-leap-smear-refid-00

Abstract

RFC 5905 [RFC5905] and earlier versions of NTP are the overwhelming method of distributing time on networks. Leap Seconds will continue to exist for a good number of years' time, and since the timescale mandated by POSIX effectively ignores any instances where there are not 86,400 seconds' time in a day, something must be done to reliably synchronize clocks during the application of leap second corrections. One mechanism for dealing with the application that has recently become visible is to apply the leap second using a "smear", where the time reported by leap-second aware servers is gradually adjusted so there is no major disruption to time synchronization when processing a leap second.

While leap second processing can be expected to be properly handled by up-to-date software and by time servers, there are large numbers of out-of-date software installations and client systems that are just not able to properly handle a leap second correction.

This proposal offers a way for a system to generate a REFID that indicates that the time being supplied in the NTP packet already contains an amount of leap smear correction, and what that amount is.

Table of Contents

1. Introduction 2
1.1. Requirements Language 3
2. Leap Smear REFID 3
3. Questions 4
4. Acknowledgements 4
5. IANA Considerations 5
6. Security Considerations 5
7. Normative References 5
Author's Address 5

1. Introduction

RFC 5905 [RFC5905] and earlier versions of NTP are the overwhelming method of distributing time on networks. Leap Seconds will continue to exist for a good number of years' time, and since the timescale mandated by POSIX effectively ignores any instances where there are not 86,400 seconds' time in a day, something must be done to reliably synchronize clocks during the application of leap second corrections. One mechanism for dealing with the application that has recently become visible is to apply the leap second using a "smear", where the time reported by leap-second aware servers is gradually adjusted so there is no major disruption to time synchronization when processing a leap second.

While leap second processing can be expected to be properly handled by up-to-date software and by time servers, there are large numbers of out-of-date software installations and client systems that are just not able to properly handle a leap second correction.

This proposal offers a way for a system to generate a REFID that indicates that the time being supplied in the NTP packet already contains an amount of leap smear correction, and what that amount is.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Leap Smear REFID

RFC 5905 [RFC5905] defines the data type of NTP time values in Section 6, "Data Types":

> All NTP time values are represented in twos-complement format, with bits numbered in big-endian (as described in Appendix A of [RFC0791]) fashion from zero starting at the left, or high-order, position. ...

The 32 bit signed integer seconds portion and the 32 bit unsigned fractional seconds portion, or 32:32 format is:

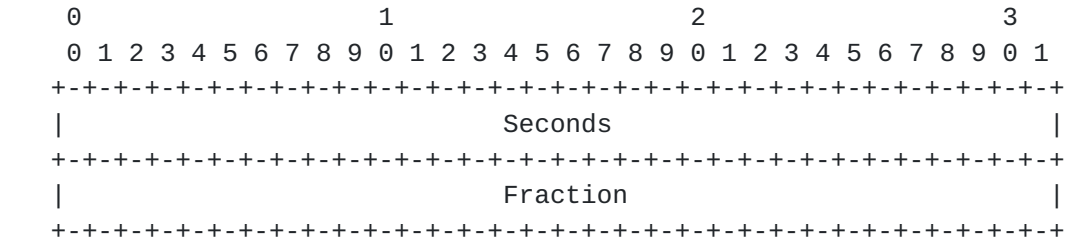

NTP Timestamp Format (32:32)

This format provides coverage for 136 years' time to a precision of 232 picoseconds. If a leap-second addition is being completely smeared just before before the stroke of the next POSIX second then the smear correction will be (0,1). If this was the only way to apply a leap smear correction then we could simply use an unsigned value to represent the correction. But while the first popular leap smear implementation applied the correction over an appropriate number of hours' time before the actual leap second so the system time was corrected at the stroke of 00:00, that meant that the difference between system time and UTC spent half of the duration of the smear application at [.5,1) "off" of correct time. The second popular implementation of the leap smear applied the first half-second correction before the stroke of 00:00 for a correction range of (0,.5] and the last half-second correction starting at the stroke

of 00:00 for a [-.5,0) correction range. This also means we need a signed value to represent the amount of correction.

The REFID of a system that is supplying smeared time to client requests while leap-smear correction is active would be 254.b1.b2.b3, where the three octets (b1, b2, and b3) are a 2:22 formatted value, yielding precision to 238 nanoseconds, or about a quarter of a microsecond.

Note that if an NTP server decides to offer smeared time corrections to clients, it SHOULD only offer this time in response to CLIENT time requests. An NTP server that is offering smeared time SHOULD NOT send smeared time in any peer exchanges. Also, CLIENT machines SHOULD not be distributing time (smeared or otherwise) to other systems.

We also note that during the application of a leap smear, the REFID from a system offering smeared time cannot provide detection of a timing loop. This is not expected to be a problem because time server systems are not expected to make CLIENT connections with each other, so they should not be receiving smeared time. Moreso, if a time server is configured to make CLIENT connections to a server that offers smeared time, with the mechanism described here it can detect when it is getting smeared time, and either ignore time from that source, or "undo" the leap smear correction and use the corrected time for that sample.

This proposal is not an attempt to justify servers offering leap smeared time. It is only an attempt to make it easy and visible to identify when a client is receiving smeared time, and provide the client to know the amount of smear correction as of the latest successful poll.

3. Questions

Should we ask IANA to allocate a pseudo Extension Field Type of 0xFFFE (for example) so the proposed "I-Do" exchange can report whether or not this server will offer leap smeared time in response to CLIENT time requests, identifying the amount of correction using the above REFID?

4. Acknowledgements

The author wishes to acknowledge Dan Mahoney (and perhaps others) for suggesting the idea of using an "impossible" first-octet value to indicate an IPv6 refid hash. The author wishes to acknowledge the contributions of Joey Saccadonuts.

5. IANA Considerations

This memo makes no requests of IANA.

6. Security Considerations

Additional information TBD

7. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <http://www.rfc-editor.org/info/rfc5905>.

Author's Address

Harlan Stenn
Network Time Foundation
P.O. Box 918
Talent, OR 97540
US

Email: stenn@nwtime.org