One-way Delay Measurement Using NTP

Vladimír Smotlacha, CESNET, Prague, Czech Republic
vs@cesnet.cz

I. INTRODUCTION

This paper presents an experimental comparison of different NTP synchronization strategies to measure one-way delay (OWD) in large area networks. Presented experimental results allow us to evaluate the accuracy of such OWD measurement.

One-way delay is one of the QoS parameters. It is defined both by the IETF as the OWD for IPPM and by the ITU-T as the IPTD (IP Packet Transfer Delay). In both cases it belongs to time-sensitive parameters, where time synchronization of both sender and receiver is required. Any error of the time synchronization affects the accuracy of the OWD measurement. One should notice that the accuracy of measurement is based on a relative synchronization and therefore an absolute synchronization to UTC is not required. Different synchronization techniques can be used. The most accurate one is to provide each site of measurement with a precise external clock. An atomic clock or a GPS receiver can be used but this solution suffers from bad scalability. The cost of an atomic clock is very high and installation of a GPS antenna may be difficult in some localities. Another method is a synchronization via network which is perfectly scalable but its accuracy is not as good.

The paper presents the direct method of one-way delay (OWD) measurement using NTP synchronization, and conditions under which the estimated error of measurement is not worse than 1 millisecond will be discussed. The measurement sites are synchronized using NTP servers. No other source of exact time is needed for the measurement, however we used a PPS (Pulse Per Second) signal from GPS to evaluate the exact accuracy of described method and to compare different setups.

All tests were done between boxes located in CESNET (NREN of The Czech Republic) and Heanet (NREN of the Irish Republic).

II. MEASUREMENT METHOD

Raw OWD was measured using the RUDE/CRUDE tool. RUDE transmits a stream of UDP packets with defined traffic shape. Each packet contains a sequential number and a timestamp of transmission time (given by the sender clock). CRUDE collects this traffic, adds timestamps of reception (given by the receiver clock) and generates a log. One can evaluate several QoS parameters from this log, e.g., throughput, loss ratio, one-way delay, and one-way delay variation.

Each measurement site runs an NTP daemon which synchronizes local clock to the NTP server. We used the parameter 'maxpoll = 6' forcing the NTP process to contact NTP server at a period not exceeding 64 s. The poll period could reach 1024 s without this parameter and we proved that the accuracy of measurement would be degraded in this case.

As a closed loopback is the base of the NTP process, the actual time offset of local clock is known. This offset is reported by the command 'ntpq -c rl' or returned by the function ntp_adjtime().

In order to know the exact local clock offset (independent on the NTP daemon), we provide each box with a PPS signal. Thus we obtain offset with an absolute accuracy of about 10 microseconds. It should be noted that the PPS signal was not used for synchronization but only for exact offset evaluation.

Let us assign:

- T_s - timestamp of sending the packet (from CRUDE log)
- T_r - timestamp of receiving the packet (from CRUDE log)
- O_s - sender clock offset (reported by NTP)
- O_r - receiver clock offset (reported by NTP)
- P_s - exact sender clock offset (from PPS capture log)
- P_r - exact receiver clock offset (from PPS capture log)

From these data we can calculate:

- Raw one-way delay obtained from CRUDE log
  $$OWD_r = T_r - T_s$$
- One-way delay corrected by NTP offsets
  $$OWD_n = T_r - O_r - (T_s - O_s)$$
- Exact one-way delay calculated from GPS time
  $$OWD_e = T_r - P_r - (T_s - P_s)$$

III. MEASUREMENT RESULTS

We performed the measurement in three setups with different location of NTP servers. In Setup I we used two NTP servers (tik.cesnet.cz and Karlovy.hea.net alias tt35.ripe.net) located in the measurement PoPs with round-trip delay less then one millisecond from each box. In Setup II we used one common NTP server (worf.ijs.si) located in Slovenia, i.e. in the GEANT network but outside both PoPs. In Setup III we used one common NTP server located in CESNET.

The CESNET measurement box was a DELL 1400 with Pentium III/860MHz. The Heanet box was a standard PC with Pentium III/450MHz. Both boxes ran Linux, kernel 2.4.16 with nanokernel patch. We used NTP version 4.1.71 and RUDE version 0.50.

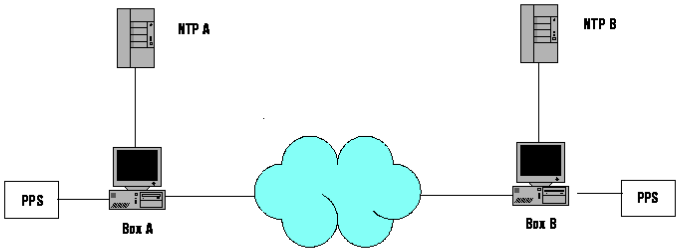

Fig. 1. Setup I

In Setup I, the round-trip delay between each measurement site and its local NTP server was less than 1 millisecond. We proved that the error of OWD measurement was in interval +-500 microseconds using Setup I.

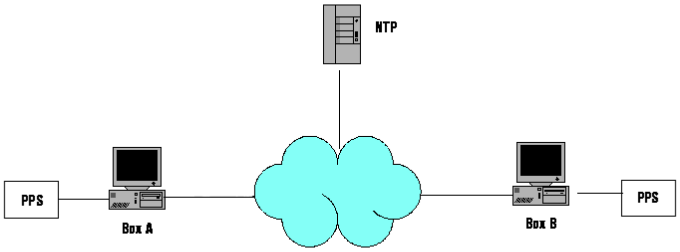

Fig. 2. Setup II

In Setup II, a single common NTP server located in a third network was used - the box worf.ijs.si in Slovenia. An average round-trip delay was 30 ms between this NTP server and CESNET, and 54 ms between this NTP server and Heanet. Therefore, conditions for good time synchronization were much worse compared to those in Setup I. We expected much worse results than in Setup I, but the error of OWD measurement was still far below +- 1 millisecond.

Setup IIa was the same as Setup II with a single difference: the parameter ’maxpoll = 6’ was omitted from the NTP daemon configuration. We got worse results than in Setup II: the error of OWD measurement was in interval -3ms - +1ms. It is evident that the polling interval parameter of the NTP daemon has to be kept small.

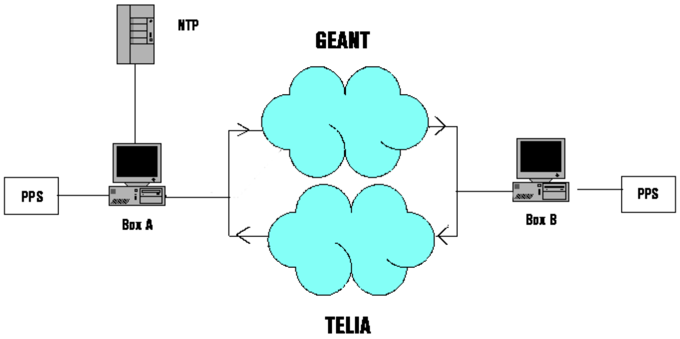

Fig. 3. Setup III

In Setup III we tested the influence of asymmetric routing. While traffic from Heanet to CESNET was routed in a standard way, i.e. through the pan-european GEANT network (OWD of about 20 ms), traffic in the opposite direction was routed via the Telia network which supplies the commodity Internet traffic to CESNET (OWD of about 37 ms). To demonstrate the influence of asymmetric routing, we used one common NTP server located in Heanet. The average error of OWD measurement was 8 ms, about one half of the OWD difference in both directions.

IV. CONCLUSION

We proved that the OWD can be measured by boxes whose clock is synchronized using the NTP protocol. The maximum measurement error is below 1 ms if an NTP server is located in each measurement site. Usage of one common NTP server offers still quite a good accuracy if we can provide symmetric routing between the NTP server and the measurement boxes.

REFERENCES

[1] D.L.Mills. “Clock Discipline Algorithm for the Network Time Protocol Version 4” Electrical Engineering Department Report 97-3-3, University of Delaware, March 1997.

[2] D.L.Mills. “Adaptive Hybrid Clock Discipline Algorithm for the Network Time Protocol” IEEE/ACM Trans. Networking 5, 6 (October 1998).

[3] D.L.Mills, P.-H.Kamp. “The nanokernel” Proceedings, Precision Time and Time Interval (PTTI) Applications and Planning Meeting (Reston VA, November 2000). of Services in IP and ATM, videoconferencing systems and communications protocols.

[4] D.L.Mills. “Network Time Protocol Specification” RFC-1305, IETF, March 1992.

[5] V.Paxson, G.Almes, J.Mahdavi, M.Mathis. “Framework for IP Performance Metrics” RFC-2330, IETF, May 1998.

[6] G.Almes. S.Kalidindi, M.Zekauskas. “A One-way Delay Metric for IPPM” RFC-2679, IETF, September 1999.

[7] J.Mogul, D.L.Mills, J.Brittenson, J.Stone and U.Windl. “Pulse-per-second API for Unix-like operating systems, version 1” RFC-2783, Internet Engineering Task Force, March 2000.