Survivable, Real Time Network Services

Defense Advanced Research Projects Agency
Contract F30602-98-1-0225, DARPA Order G409/J175

Quarterly Progress Report
1 April 2001 - 30 June 2001

David L. Mills
Electrical Engineering Department
University of Delaware

1. Introduction

This report covers the work done in support of the DARPA Information Technology Office program in computer networking. Contributors to this effort include Prof. David L. Mills and graduate Tamal Basu. Graduate student Qiong Li has completed his dissertation and been granted the PhD degree. The project continues previous research in network time synchronization technology jointly funded by DARPA and US Navy. The technology makes use of the Network Time Protocol (NTP), widely used in the Internet, together with engineered modifications designed to improve accuracy in high speed networks. Specific applications benefiting from this research include multicast topologies, multimedia, real-time conferencing, cryptographic systems, and management of distributed, real-time systems.

This quarterly report is submitted in traditional report form on paper. As the transition to web-based information dissemination of research results continues, almost all status information and progress reporting is now on the web, either on pages belonging to the principal investigator or to his students. Accordingly, this and future progress reports will contain primarily schedule and milestone data; current status and research results are reported on web pages at www.eecis.udel.edu/~mills in the form of papers, technical reports and specific briefings.

2. Autonomous Authentication

The IETF STIME group has reviewed the revised Internet Draft submitted some months ago and provided a list of suggested changes and additions. Some concern was expressed about the cookie used in client/server mode; in particular, the relative ease an intruder can intercept it and use it to disrupt protocol operations. However, the cookie has only limited usefulness, since a passive intruder cannot predict the next session ID and cannot synthesize a new message believable to the client. However, an active intruder could intercept the server reply and use the cookie to forge messages that would be believed. Even in this case there is a good chance that one or more of the eleven NTP sanity checks would deflect such attacks. Notwithstanding these caveats, the reviewers agreed the need for some way to prevent incidental disclosure of the cookie.

To deal with this problem the Autokey protocol client/server mode operations were modified to include a Diffie-Hellman agreement exchange where the agreed key is used as a stream cipher to encrypt and decrypt the cookie. The agreed key is discarded after use, so the server remains stateless. The scheme has been implemented and tested in the Autokey Version 2 protocol.

The STIME reviewers suggested additions to the protocol that would support certificate transfer directly by the protocol. Formerly, it was assumed that certificates would be obtained and verified using other means using generic protocols for just that purpose. However, it was agreed that provisions for inband transport would be a useful addition to the protocol. This has been done in the Autokey Version 2 protocol, where certificates, Diffie-Hellman parameters and leapsecond tables are now automatically transported from servers to clients via the protocol.

There has been no positive response from the RSA Laboratories for permission to include the rsaref20 cryptographic software library with the NTP distribution. While the RSA patents have run out, that software is copyrighted; however, users have been able to find sources for compatible software in Europe and elsewhere. This situation points up the need to find open software with equivalent functionality from some other place. A suitable source has been found in the OpenSSL cryptographic library, which includes in addition to the cryptographic routines provisions to construct X.509 certificates and certificate requests, sign certificates and in general implement the functions of a certificate authority.

We have completed the integration of the OpenSSL software in the NTP distribution. This has resulted in an extensive redesign of the Autokey protocol, which has resulted in the Autokey Version 2 described on web pages and briefings and in the NTP software documentation. The new version uses only one packet format, requires fewer messages and is more robust to intrusion and packet loss. The primary cryptographic defense is the X.509 certificate, which is compatible in format and content with industry standard practices. In addition, the full complement of cryptographic message digest and signature encryption routines in the OpenSSL library are available.

The current status of the Autokey protocol and implementation is available via the web at www.eecis.udel.edu/~mills/autokey.htm.

3. Autonomous Configuration

Work continues on the development and test of the Manycast autonomous configuration scheme. In this scheme a multicast client sends an ordinary request packet to a multicast group address and potential multicast servers in TTL range have the opportunity to respond with an ordinary unicast message. Each received message causes the client to mobilize an ephemeral association, which then proceeds as if the server were originally explicitly configured.

Initial experiments with the DCnet routers and CAIRN routers have demonstrated a clique of Autokey/Manycast routers do in fact automatically reconfigure when something breaks and recover the necessary security configuration as required. However, experience has shown an unexpected incidence of “clockhopping” where many more than three potential sources are avaialable. What happens is that the clustering algorithm attempts to sift the best three sources out of a possibly large population, but several combinations of three sources exist which have substantially the same performance metric. It is likely in such cases that the particular combination of three sources found in succeeding messages is different, so that one or more of the current survivors are ejected in favor of a new, substantially equivalent one. As each new association requires a new proventication exchange, the result is an necessary protocol exchange and processor burden.

The current status of the Manycast mode and implementation is available via the web at www.eecis.udel.edu/~mills/autocfg.htm.

4. The Huff-n'-Puff Filter

A particularly interesting feature has been developed for the NTP protocol that deserves a section all its own. In cases involving considerable network congestion the NTP clock filter and clustering algorithms do a decent job of casting off statistical outlyers. However, it has long been recognized that there is some degree of additional statistical information that is not being utilized. As described in the web briefings and previous papers, the distribution of samples on a graph plotting offset against delay appears as a wedge with apex at the minimum delays found on the path. Typically, as congestion increases along the path the samples at larger delays are distributed within the wedge between limbs extending from the apex at positive and negative slopes of 0.5.

A tantalizing observation has been the characteristic as congestion increases is the tendency for samples to first lie near the limbs and then fill in toward the center of the wedge. Thus, if a way could be discovered to select just the samples near the limbs, these sample values could be corrected relative to the apex and used to accurately discipline the system clock.

The problem under conditions of high congestions is that there is no accurate way to estimate the minimum delay and position the apex of the wedge diagram. After some discussion with others of the NTP developers group and considerable simulation using Matlab, a scheme called Huff-n'-Puff (simply huffpuff hence) has been developed and implemented in the NTP protocol. The scheme is particularly useful for sites that have a relatively low speed connection to the Internet and spend considerable time during part of the day either downloading or uploading large files. Under these conditions there can be large delays on one direction of transmission, but much smaller delays on the other.

Under these conditions the wedge diagram samples cluster along one of the limbs and are easily corrected if the minimum roundtrip delay is known. For this purpose, the huffpuff scheme includes a minimum filter which selects the minimum delay sample in a window extending from minutes to hours, depending on the anticipated maximum duration of congestion over the day. During congested periods delay samples above or below the nominal offset of the wedge are corrected before use as discipline source.

As confirmed by reports from users where formerly the NTP timekeeping was so bad that it had been turned off during the day, the performance of huffpuff has made a dramatic improvement to the point that NTP can now be run continuously. However, there still remains room for improvement in cases where the congestive delays are distributed more evenly over both directions of transmission. The best defense implemented at the moment is a popcorn spike filter that discards large outlyers from the limbs. It appears an approach modelled on an expert-system might be more effective at capturing the best samples in such cases.

5. The NTP Book

Work on the NTP book continues. Eleven chapters are now substantially complete, although considerable work remains on figures, tables and cross references. Originally, this was to be an O'Reilly book with two additional authors. This investigator was to contribute the technical

description, analysis and experiment chapters and the other two authors were to contribute hands-on programming manual and anecdotal experience in the O'Reilly tradition. As the work progressed, one of the other authors had a serious automobile accident and was unable to complete his assignment. The third author also failed to complete his assignment. This investigator concluded that the eleven chapters are too technical for an O'Reilly book anyway, so opted out of the O'Reilly contract. A suitable publisher has yet to be identified.

6. Infrastructure

There have been a continuous stream of contributions to the NTP software distribution from the developer's group. These include further refinements to the clock filter and clock discipline algorithm, new reference clock drivers and the huff-n'-puff filter mentioned above. The latest version of NTP Version 4 has been deployed in the DCNet research network, EECIS campus network and CAIRN research network.

The web documentation for this investigator's home page, the DCnet and CAIRN pages, the NTP home page and the NTP documentation pages have all been brought up to date. These pages have proliferated to the better part of a hundred megabytes with text, pictures and sounds. The web should be much easier to walk and the pages much more interconsistent and accurate. Besides all papers, technical reports, technical memoranda published at Delaware for the last 14 years, there are a number of technical briefings on projects funded by DARPA, US Navy and US Army agencies. DARPA status reports, yearly reports, final reports, quad charts and news stories are all archived at www.eecis.udel.edu/~mills/support.htm. Desktop wallpaper and a couple of interesting historic lessons are included in the www.eecis.udel.edu/~mills/gallery.htm.

7. Future Plans

Our plans for the next quarter include further work on the NTP book mentioned above. Additional documentation plans include upgrading the web pages on the simulation of very large networks to include material from Mr. Basu's masters thesis.

We plan to continue testing and refining the Autokey Version 2 protocol and implementation. Specific goals are to integrate provisions to generate Digital Signature Standard (DSS) certificates in the key generation routines. We also plan to upgrade to version 3 of the X.509 certificate and include provisions for automatically verifying the certificate trail. Eventual plans are to integrate a certificate discovery function in the manycast scheme, which is necessary for a field-deployable sensor network.

We plan to continue work on the Autoconfigure scheme, in particular to develop means to avoid the clockhopping problem mentioned earlier. Other improvements expected include a rewrite of the configuration code and name resolution code.

8. Publications

All publications, including journal articles, symposium papers, technical reports and memoranda are now on the web at www.eecis.udel.edu/~mills. Links to the several publication lists are available on that page, as well as links to all project descriptions, status reports and briefings. All publications are available in PostScript and PDF formats. Briefings are available in HTML,

PostScript, PDF and PowerPoint. The project descriptions are cross-indexed so that the various interrelationships are clearly evident. Also included are the documentation pages for various public software distributions. Links to other related projects at Delaware and elsewhere are also included on the various pages. Hopefully, the organization of these pages, which amount to a total of about 300 megabytes of information pages and reference documents, will allow quick access to the latest results and project status in a timely way.

Following is a retrospective list of papers and reports supported wholly or in part on this project and the immediately preceding project "Scalable, High Speed, Internet Time Synchronization," DARPA Order D012. The complete text of all papers and reports, as well as project briefings, status reports and supporting materials is at www.eecis.udel.edu/~mills.

8.1 Papers

1. Levine, J., and D. Mills. Using the Network Time Protocol to transmit International Atomic Time (TAI). *Proc. Precision Time and Time Interval (PTTI) Applications and Planning Meeting* (Reston VA, November 2000).

2. Mills, D.L., and P.-H. Kamp. The nanokernel. *Proc. Precision Time and Time Interval (PTTI) Applications and Planning Meeting* (Reston VA, November 2000).

3. Mills, D.L. A brief history of NTP time: confessions of an Internet timekeeper. Submitted for publication; please do not cite or redistribute.

4. Li, Qiong, and D.L. Mills. Investigating the scaling behavior, crossover and anti-persistence of internet packet delay dynamics. *Proc. 1999 IEEE GLOBECOM 99 Symposium* (Rio de Janerio, Brazil, December 1999).

5. Mills, D.L. Cryptographic authentication for real-time network protocols. In: *AMS DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 45* (1999), 135-144.

6. Mills, D.L. Adaptive hybrid clock discipline algorithm for the Network Time Protocol. *IEEE/ACM Trans. Networking 6, 5* (October 1998), 505-514.

7. Mills, D.L., T.S. Glassey, and M.E. McNeil. Authentication scheme extensions to NTP. Internet Draft draft-mills-ntp-auth-coexist-01.txt, University of Delaware and Glassey-McNeil Technologies, September 1998, 9 pp.

8. Li, Qiong, and D.L. Mills. On the long-range dependence of packet round-trip delays in Internet. *Proc. IEEE International Conference on Communications* (Atlanta GA, June 1998), 1185-1191.

9. Mills, D.L., A. Thyagarajan and B.C. Huffman. Internet timekeeping around the globe. *Proc. Precision Time and Time Interval (PTTI) Applications and Planning Meeting* (Long Beach CA, December 1997), 365-371.

10. Mills, D.L. Authentication scheme for distributed, ubiquitous, real-time protocols. *Proc. Advanced Telecommunications/Information Distribution Research Program (ATIRP) Conference* (College Park MD, January 1997), 293-298.

11. Mills, D.L. The network computer as precision timekeeper. *Proc. Precision Time and Time Interval (PTTI) Applications and Planning Meeting* (Reston VA, December 1996), 96-108.

12. Mills, D.L. Improved algorithms for synchronizing computer network clocks. *IEEE/ACM Trans. Networks 3, 3* (June 1995), 245-254.

8.2 Technical Reports

13. Mills, D.L. Public key cryptography for the Network Time Protocol. Electrical Engineering Report 00-5-1, University of Delaware, May 2000. 23 pp.

14. Mogul, J., D. Mills, J. Brittenson, J. Stone and U. Windl. Pulse-per-second API for Unix-like operating systems, version 1. Request for Comments RFC-2783, Internet Engineering Task Force, March 2000, 31 pp.

15. Sethi, A.S., H. Gao, and D.L. Mills. Management of the Network Time Protocol (NTP) with SNMP. Computer and Information Sciences Report 98-09, University of Delaware, November 1997, 32 pp.

16. Mills, D.L. A precision radio clock for WWV transmissions. Electrical Engineering Report 97-8-1, University of Delaware, August 1997, 25 pp.

17. Mills, D.L. Clock discipline algorithms for the Network Time Protocol Version 4. Electrical Engineering Report 97-3-3, University of Delaware, March 1997, 35 pp.

18. Mills, D.L. Proposed authentication enhancements for the Network Time Protocol version 4. Electrical Engineering Report 96-10-3, University of Delaware, October 1996, 36 pp.

19. Mills, D.L. Simple network time protocol (SNTP) version 4 for IPv4, IPv6 and OSI. Network Working Group Report RFC-2030, University of Delaware, October 1996, 18 pp.

20. Mills, D.L. Simple Network Time Protocol (SNTP). Network Working Group Report RFC-1769, University of Delaware, March 1995, 14 pp.

21. Mills, D.L. Simple Network Time Protocol (SNTP). Network Working Group Report RFC-1769, University of Delaware, March 1995, 14 pp.

8.3 Internet Drafts

22. Mills, D.L. Public-Key Cryptography for the Network Time Protocol. Internet Draft draft-ietf-stime-ntpauth-01.txt, University of Delaware, April 2001, 45 pp.

23. Mogul, J., D. Mills, J. Brittenson, J. Stone and U. Windl. Pulse-Per-Second API for UNIX-like Operating Systems, Version 1.0. Internet Draft draft-mogul-pps-api-05.txt, Compaq Western Research Laboratory, August 1999, 30 pp. (obsoleted by RFC-2783)

24. Mills, D.L., T.S. Glassey, and M.E. McNeil. Authentication Scheme Extensions to NTP. Internet Draft draft-mills-ntp-auth-coexist-01.txt, University of Delaware and Glassey-McNeil Technologies, September 1998, 9 pp. (expired)