Scalable, High Speed, Internet Time Synchronization

Advanced Research Projects Agency
Contract DABT 63-95-C-0046

Quarterly Progress Report
1 June 1996 - 31 August 1996

David L. Mills
Electrical Engineering Department
University of Delaware

1. Introduction

This report covers the work done in support of the Information Technology Office of the DARPA on reliable, accurate network time synchronization. Contributors to this effort include Prof. David L. Mills and graduate students Ajit Thyagarajan and Bradley Cain.

The project continues previous research in network time synchronization technology jointly funded by US Navy, US Army, DARPA and NSF. 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, including SONET and ATM, expected to be widely deployed in the next several years. Specific applications benefiting from this research include multicast topologies, multimedia, real-time conferencing, cryptographic systems, and management of distributed, real-time experiments.

Recent projects reported in papers, technical reports, project reports and technical memoranda include precision timekeeping devices for Sun Microsystems workstations, improved clock discipline algorithms, and protocol upgrades for the Network Time Protocol (NTP). Other projects include a relatively inexpensive precision timing receiver using the LORAN-C radionavigation system, and an optimum matched-filter receiver/decoder using DSP technology. Software developed with joint funding includes the NTP Version 3 implementation for Unix and Windows and a set of precision-time kernel modifications for major Unix workstation manufacturers. Finally, the joint projects involve the conduct of experiments designed to evaluate the success of the research and assist technology transfer to computer manufacturers and network providers.

2. Present Status

Most of the effort during the quarter were concentrated on the design and documentation of the Network Time Protocol, Version 4. The results of an extensive survey of NTP servers and clients are reported. In addition, several hardware and software enhancements to our growing zoo of workstations, routers and infrastructure support were implemented. Two new developments are reported, one involving a series of intrusions due to apparent software configuration errors elsewhere in the Internet, and another involving an experiment in long-range dependency phenomena in the global timekeeping subnet.

2.1 Publications

Two technical reports and one Internet draft were published during the quarter. Two abstracts were accepted for papers to be presented at future conferences.

2.1.1 Technical Reports

The following reports were produced during the quarter:N

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, 14 pp. Also published in PostScript: *Ibid*. Electrical Engineering Report 96-10-2, University of Delaware, October 1996, 14 pp. URL: see http://www.eecis.udel.edu/~mills/reports.html.

Abstract

This report describes the Simple Network Time Protocol (SNTP) Version 4, 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. This report obsoletes RFC-1769, which describes SNTP Version 3. Its purpose is to correct certain inconsistencies in the previous document and to clarify header formats and protocol operations for current NTP Version 3 (IPv4) and proposed NTP Version 4 (IPv6 and OSI), which are also used for SNTP.

Mills, D.L. Proposed Authentication Enhancements for the Network Time Protocol Version 4. Electrical Engineering Report 96-10-3, University of Delaware, October 1996, 37 pp. URL: see http://www.eecis.udel.edu/~mills/reports.html.

Abstract

This report describes proposed changes in the security model and authentication scheme for the Network Time Protocol Version 4, which is an enhanced version of the current Versing 3. The changes are intended to replace the need to securely distribute cryptographic keys in advance, while protecting against replay and man-in-the-middle attacks. As in other schemes described in the literature, the proposed scheme is based on the use of a public-key cryptosystem to verify a server secret and from this to generate session keys for each client separately. A particularly important consequence of this design in the case of NTP is that the mechanisms for time synchronization and cryptographic signature verification must be decoupled to preserving good timekeeping quality. The schemes to do this are the main body of this report, which also includes an extensive analysis of the vulnerabilities to various kinds of hardware and software failures, as well as hostile attack.

This report has been submitted as an RFC in ASCII format to the RFC Editor. As there are a number of figures in the document which cannot be rendered in ASCII, the above document in PostScript version must be considered the definitive one and would ordinarily be published as an RFC in PostScript format. However, current procedures require RFCs in PostScript format must be identical to the ASCII version, except for minor formatting differences. Accordingly, the RFC submission is in ASCII format with no mention of the figures and explanatory text. There are no plans to submit this RFC in PostScript format.

Besides the specific proposals for NTP Version 4, the report contains an intricate vulnerability analysis of the NTP security model and authentication scheme. It is expected that portions of this report will be developed as papers to be submitted later.

2.2 Precise Time and Time Interval (PTTI) Symposium

An abstract for a paper on precision computer network time synchronization has been accepted for presentation at the PTTI Symposium to be held in Washington, DC, in December. Following is the text of the abstract:

> General purpose workstation computers are becoming faster each year, with processor clocks now operating at 300 MHz and faster. Computer networks are becoming faster, with speeds of 155 Mbps available now and 2.4 GHz being installed. Using available technology and existing workstations and Internet paths, it has been demonstrated that reliable computer network time synchronization can be achieved with accuracies better than a millisecond in LANs and better than a few tens of milliseconds in many places of the global Internet. This technology includes the Network Time Protocol (NTP), now used in an estimated total of over 100,000 clients and servers in a distributed, hierarchical, synchronization network spanning the Internet. Over 100 primary time servers are at the roots of this network, each connected to an external source of time, such as a GPS receiver or telephone modem.
>
> Reliable network synchronization requires crafted algorithms which minimize jitter on diversity paths between servers and clients, determine the best subset in a set of redundant servers, and discipline the computer clock in both time and frequency. This paper describes a number of these algorithms with particular emphasis on achieving the best accuracy possible with existing hardware and operating systems software. With these tools and unmodified workstations and LANs, it is possible to discipline network computer clocks to less than a millisecond in time and less than a few parts per million in frequency.
>
> However, the true adventure is to develop refined algorithms and, where necessary, hardware which can directly discipline the computer clock to an external source, such as a pulse-per-second signal or externally disciplined counter peripheral. This technique has been used to remove jitter accumulated on long paths across the Internet, while using NTP to provide UTC seconds numbering. This paper describes hardware and software modifications to several families of Unix workstations which provide time synchronized to within a few microseconds of UTC and frequency stabilized to within a fractional part per million without hardware modifications. This requires Unix operating system kernel modifications which implement a software-disciplined, hybrid phase/frequency-locked loop totally contained in the kernel. These modifications have been incorporated in the standard Digital Unix 4.0 operating system for the Alpha workstation.
>
> The significance of this work is confirmation that general purpose workstations with appropriate software modifications can be used for critical time measurements with accuracies in the tens of microseconds. This may be specially useful for astronomy, oceanography, manufacturing and process control applications.

2.2.1 Army Research Laboratories (ARL) Conference

An extended abstract for a paper on NTP security issues has been accepted for presentation at the ARL Conference to be held in College Park, MD, in January 1997. Following is the text of that abstract:

> Introduction
>
> The Network Time Protocol (NTP) is widely used in the Internet to synchronize computer time to national standards. The current NTP subnet population includes well over 200 primary (stratum-1) servers and 100,000 secondary (stratum-2 and above) servers and clients with over 500,000 continuously operating protocol instantiations between them. NTP 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. The protocol uses redundant servers, diverse network paths and crafted algorithms which cast out incorrect servers and minimize errors due to network delay jitter. 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.
>
> A reliable, ubiquitous and efficient time synchronization service requires some provision to prevent accidental or malicious attacks on the servers and clients of the service. Reliability requires that clients can determine that received messages are authentic; that is, were actually sent by the intended server and not manufactured or modified by an intruder. Ubiquity requires that any client can verify the authenticity of any server using only public information. Efficiency requires that the resources required by the protocol instantiations cause minimal performance impairment to the other services supported by the hardware and software.
>
> A public-key cryptosystem such as RSA [PKS93a] would be a good foundation on which to build the authentication scheme. Public-key cryptography has the very valuable characteristic that the server and client need share no common variables other than those declared public values. In order to minimize the vulnerability to cryptographic attack, every message must be individually signed using the server private key. However, the computational burden to this is unacceptable on a busy server with many hundreds of clients. A key-agreement scheme such as Photuris [KAR95], SKIP [AZI95] or ISAKMP [MAU95] can be used to construct a shared secret used as a session key between two peers. However, session keys specific to each client are impractical in servers supporting many hundreds of clients. This paper describes crafted algorithms which provide reliable authentication while avoiding these penalties.
>
> Requirements
>
> An important requirement for the authentication scheme is that it provide for automatic discovery of authenticated servers whose identities (IP addresses) are not known in advance. In NTP clients can discover servers directly using multicast advertisements sent by gratuitous (and possibly compromised) servers. In addition, servers can be discovered by their responses to multicast invitations sent by clients. It is necessary to authenticate each discovered server using public values cryptographically bound to that server with signed certificates. The certificates themselves cannot be considered reliable until their lifetimes are verified. In principle, lifetimes cannot be trusted, unless all clocks in the certificate network

are synchronized. This requires that the time synchronization function and the certification function operate independently on a tentative basis and be prepared to retry operations that fail due to indeterminate outcomes.

The security model specifically recognizes that synchronization and authentication services may not be continuously available, due to attack or failure. Peers can fail or operate incorrectly or an intruder can attempt to modify messages or jam the subnet in one way or another. Transmission lines can fail, routes can change or become congested, and cryptographic keys and even security policies can change while the subnet is in regular, continuous operation. Thus, the design approach must use redundant resources and diverse communication paths. And provide for an exceptional degree of fault resistance.

Design Approach

The three modes of NTP operation: peer/peer, client/server and multicast/anycast present quite different security models. In peer/peer modes, both peer associations maintain persistent state variables, so conventional methods such as Photuris are acceptable. In client/server modes, the server session key is regenerated for each received request and must be different for each client. First, the client sends a request to the server to compute the session key, which is a hash of private values and the client address. The hash is encrypted and returned to the client. This method prevents a man-in-the-middle attack, but does consume significant server and client resources. However, the exchange needs to be done only infrequently when the client is first started up and when the server private values are changed. The session key is cached by the client and used to compute the message digest. The server recomputes the session key as each request is received.

In multicast mode, a server first calculates a list of session keys for later use, as in the S/KEY system [HAL95]. The first entry in the list is determined as above. The second entry is calculated as a hash of the first entry and certain public values. Continuing in this way, the server fills the list, which may have from a few to several hundred entries. The server uses the list in inverse order; that is, the last entry is used first, then the next before that, and so on until all entries have been used. At this point, the server generates a new list. A client authenticates each message relative to the immediately preceding message. It first obtains the current session key, which is a public value, and RSA signature from the server. The client accepts following messages only if each one hashes to the immediately preceding session key. If not, a message might have been lost in transit, so the client hashes again. This procedure may continue for a fixed number of hashes, following which the client abandons the attempt and obtains the current session key from the server.

2.3 Software and Hardware Upgrades

The effort on software upgrades reported for the last quarter continued in the current quarter. The Hewlett Packard company provided the current version of HP-UX, Version 10.02, which is being installed on our HP 9000/735 workstation. This upgrade is significant in that the new HP-UX version includes the Unix kernel adjtime() system call, which is used by the NTP daemon to slew the local clock to a designated offset. Formerly, this had to be done by a HP-UX specific hack using a special adjtimed daemon. In addition, two of our Sun workstations were upgraded to Solaris 2.5.1 and various security related patches and new versions of the teleconferencing tools were installed.

A good deal of network reorganization has been done to accomodate a number of donated equipments, including:

1. The TrueTime company donated a time server which operates much like the 2000 TymeServe NTP server previously donated by the Bancomm company. Both devices have a self-contained GPS receiver and Ethernet interface. They provide public SNTP service as specified in RFC-1769.

2. The US Naval Observatory donated a pair of Hewlett Packard 5061A Cesium Beam Standards. Both of these have new high stability cesium tubes, which is a significant factor, since these tubes usually wear out in ten years or so and are expensive to replace. One of these devices has been installed at the campus laboratory as backup for another 5061A on indefinite loan from the US Coast Guard. The other has been installed at the backroom test site and used in connection with the Arbiter GPS receiver.

3. The Sun Microsystems company donated an UltraSPARC 170e workstation for use in IPv6 experiments, NTP protocol development and 64-bit kernel timekeeping development. When the workstation and peripherals have been integrated, a IPv6 software package provided by Sun will be installed and the machine connected to the 128.4 research net. It will be used also as a DARTnet test host for experiments with other IPv6 collaborators.

The Austron 2200A GPS receiver previously donated by the Delmarva power utility was returned from repair at the factory; however, problems with that receiver remain to be resolved. A companion Austron 2201A GPS receiver purchased some years ago has been relocated to a file server operated on behalf of the campus and department community. The current equipment and workstation configuration is documented on the web at http://www.eecis.udel.edu/~mills/lab.html.

A third Sun workstation has been equipped for teleconferencing with the purchase of a remote-controlled Cannon camera and SunView interface. These devices will operate with existing wired and wireless microphones and an audio mixer on occasions where a meeting is to be broadcast from one of our department conference rooms.

2.4 Collaboration Projects

The CAIRN collaboration has suffered a slow start, both due to the press of current projects at UCL, SAIC and our laboratory, as well as plans for the SAIC hardware connection at Washington, DC. The U Delaware near-term plans and status report has been supplied as requested by the DARTnet/CAIRN directorate. A status report and set of briefing slides is on the web at http://www.eecis.udel.edu/~mills/status.html.

We have been actively promoting an agenda to encourage other organizations to adopt a GPS radio clock and provide primary (stratum 1) time service to the community. Richard Schmidt of US Naval Observatory has installed several time servers at NRL, Washington University at Saint Louis, DEC Palo Alto Laboratories and MIT, and plans to install several more, possibly overseas. Judah Levine at NIST has installed three primary servers at NCAR and NIST in Boulder and Microsoft near Seattle. All three of these are synchronized to the NIST cesium farm using modems and the Automated Computer Time Service (ACTS). In addition to these machines and several others operated by national administrations, a new public primary server has come online

in Italy and connected to the Italian national cesium standard. We assisted in all of these installations as requested to resolve minor problems.

As found from operational experience, a third primary server is urgently needed for DARTnet. Partly as the result of some prodding from us, ISI has obtained a GPS receiver and plans to connect it to the DARTnet router there. The result should be more reliable timekeeping when either or both of the other primary servers at U Delaware and LBL come off the rails. This raises the issue of how to engineer the CAIRN timekeeping network, which is now under study.

2.5 NTP Survey

The NTP survey reported last quarter has been completed; however, provisions have been made to incrementally update the somewhat massive data base from time to time. This survey was conducted as a background process, so as not to affect traffic statistics in any measurable way. The intent is to determine reachability, time and frequency error distributions, stratum-level histograms and other related data. To date, we have found some 35,000 clients and servers running NTP on the Internet and operating over 180,000 peer-peer and client-server paths. However, this is only a fraction of the suspect population, as many sites are beyond firewalls, some use the RPC program ntpdate and some synchronize only among their own group. None of these sites are visible using the available monitoring tools and purpose-built detective kit.

The results of the survey, including breakdowns by stratum, synchronization source, stability, accuracy and error tallies, are available as briefing slides on the web http://www.eecis.udel.edu/~mills/ntp.html. The data collected and inferences made from them are to become the basis of a paper to be submitted in the next quarter.

2.6 NTP Version 4 Progress

A major upgrade to the current NTP daemon implementation for Unix and Windows is under way. The upgrade consists of a conversion to the gnu autoconfigure scheme, which greatly simplifies program maintenance and porting to new architectures. However, of the over two dozen ports of the code to widely varying architectures and operating systems, some involve quirky abuse of the traditional Unix programming environment and systems interface conventions. At the moment, ports to most common Sun, DEC, HP, SGI and BSD-derived systems is complete.

As mentioned above, changes in the SNTP Version 3 specification to support SNTP Version 4 IPv4, IPv6 and OSI have been identified and published. It is the intent that these changes also apply to NTP Version 4 specification. The reason why the SNTP specification was published before the NTP Version 4 specification is finalized is to give the community a chance to review and comment on the addressing features, which are considered very mature at this time.

Further information on NTP, including status reports and briefing slides, is available on the web at http://www.eecis.udel.edu/~mills/ntp.html. Specific projects with progress to report are as follows:

2.6.1 NTP Version 4 Autonomous Configuration

Graduate student Ajit Thyagarajan is in the final stages of design for the autonomous configuration scheme planned for NTP Version 4. Much of the work, including the intricately crafted asso-

ciation model is documented in a technical report now nearing completion. In overall plan, the scheme operates on two levels. In the first level, span-limited multicasting and anycasting paradigms are used to discover potential peer paths on the basis of the IP time-to-live (TTL) field and a concentric-ring search. In order to avoid needless multicast messages and anycast implosions with large client and server populations, carefully crafted schemes are necessary to insure only a subset of all possible servers respond to client requests and only a fraction of the clients are allowed to requests in any one polling interval. The precise modelling and simulation of these schemes is now in progress.

In the second level, the existing NTP algorithms automatically select the best subset of servers for each client separately. However, in order to minimize network traffic while maximizing accuracy, these algorithms have been modified slightly. The modified algorithms have been implemented in the current NTP daemon for Unix and Windows and are now in test. Further information on this project, including a status report and set of briefing slides, is on the web at http://www.eecis.udel.edu/~mills/autonomous.html.

2.6.2 NTP Version 4 Authentication Scheme

Graduate student Bradley Cain is developing plans to implement the proposed security model and authentication scheme described in the technical report mentioned earlier in this report. The primary impact of the authentication scheme is that clients and servers can generate cryptographic keys on-the-fly and verified by public-key algorithms, yet the resources required to perform the public-key operations do not adversely affect the accuracy of the timekeeping functions. Further information on this project, including a status report and set of briefing slides, is on the web at http://www.eecis.udel.edu/~mills/authentic.html.

2.7 Infrastructure

The massive overhaul and upgrade of personal web pages and various project web pages reported during the previous quarter continues. To claim that it has been completed, as previously reported, is simply not the case. While research status reports (not administrative status reports like this one), briefings and all papers, technical reports, project reports and technical memoranda published by our laboratory in the last ten years are on the web, maintenance of the huge amount of this stuff and updating status reports and briefing documents is a continuing burden. The most recent updates can be found in the http://www.eecis.udel.edu/~mills and http://www.eecis.udel.edu/~millslab web pages. Briefing slides for current seminar presentations and program review meetings are also available on these pages, as well as course syllabi and student lecture notes.

The web overhaul is part of paradigm shift in our infrastructure for paper and report production, documentation and communication. Retraining and transition from older Unix-based tools to newer Windows-based tools has been painful, but believed in the best interests of the teaching and research program. For paper and report production, we have abandoned Ventura Publisher and embraced FrameMaker for both Unix and Windows. For briefing slide production, we have abandoned Ventura Publisher and embraced Windows Power Point, Windows Excel and Corel Photo-Paint, as well as Unix and Windows support utilities.

2.8 Intruder Watch

During the quarter, there were three incidents involving apparently accidental intrusion on resources of the 128.4 research network. All three of these involved misconfigured hosts on other networks attempting to access nonworking addresses on the 128.4 network. All were reported to the CERT for archiving, but none required direct assistance or intervention by the CERT.

The first intrusion turned out to be an interesting exercise in network architecture and highlighted some old advice now being relearned. Some host on the Data General corporate network near Boston apparently believed itself to have address 128.4.2.1, which is not a working address on the 128.4.2 subnet. This subnet is a bridged Ethernet using an ISDN line as the connection medium. It happens this subnet is continuously monitored using tcpdump in a workstation window. A Compaq PC happens to be connected to this subnet. A LED on this machine flashes when a frame is received by the PC, including broadcast frames.

Apparently, the host wished to poll its neighbors for a game of Blackjack and sent repeated broadcasts to its purported local net broadcast address 128.4.2.255. However, a Sprint router connecting the Data General corporate net to the Sprint backbone noticed the error and sent a ICMP Unreachable message to the alleged source. That message was forwarded to the 128.4.2 subnet router, which then sent a number of repeated ARP messages to find the source. The result was a rain of ARP messages causing congestion on the ISDN line and continuous flashing of the PC LED mentioned above. This was how the intrusion was first recognized.

As past chairman of the Internet Architecture Task Force and author of the original Gateway Requirements RFC, this PI deplores the absence of wisdom in the actions of the Sprint router. Routers should never send ICMP messages to hosts whose addresses are suspected bogus. While the abuse has since stopped (the perpetrator was never found), the lessons of the past must be relearned. This is the gist of the advice tendered to all network operators involved.

The second intrusion involved a host on the TRW corporate network in California. It involved sporadic DNS queries to a nonworking address 128.4.2.1, causing much the same disruption of the 128.4.2 subnet as the first case. This turned out to be a simple typographical error; the TRW DNS server has address 129.4.2.1. The third intrusion was similar, in that a host intended to access the rackety.udel.edu time server at its 128.4.1.1 address, but had been misconfigured for nonworking address 128.4.2.1.

These intrusions were all noticed by casual observation of the PC LED and tcpdump window, not by any formal intrusion detection schemes. However, the experience points up the need for some additional monitoring of ARP traffic. Hosts should keep a list of addresses found unresponsive to ARP requests and return ICMP Unreachable (code Host Unreachable) messages to the source. In the early days of the Internet, the BBN routers in fact did this. If the abuse persists for some time, like an hour or two as in all intrusions noted above, some remedial action should be taken.

2.9 Long Range Dependency

As somewhat of a shot in the dark, graduate student Qiong Lin has been collecting reams of data representing one-way delays in the global Internet. Intrepid NTP primary time server pogo.udel.edu has been configured to watch 23 other primary time servers on all continents except Antarctica (coming soon?). The peer-peer paths have very widely varying characteristics,

some by domestic wire or fiber, some by undersea cable and some by satellite, but all are individually synchronized to UTC by radio, satellite or modem to within a millisecond or two.

The data collected represent one roundtrip (outbound and return) NTP message volley every 64 seconds, which consists of four timestamps recording the departure and arrival of both the outbound and return messages. This design allows the one-way delays on each direction to be accurately determined. Every volley for each of the 23 servers has been archived for the last three months, allowing analysis over periods from about one minute to three months. A preliminary analysis of these data using techniques developed in the literature for long-range dependencies clearly demonstrates these dependencies are revealed in the collected data. Such dependencies have been reported in the literature before, but not at the scales made possible in the present experiment design. We intend to pursue this interesting phenomena and report the results.

3. Plans for the Next Quarter

Our plans for the next quarter include continued development of the NTP Version 4 protocol model, specification and implementation. Specifically, we plan to begin implementation of the new authentication scheme and integration with the current NTP Version 3 daemon for Unix and Windows. The daemon is to be tested first in the research net, then the DARTnet/CAIRN community. As the extensions are backwards compatible, the new features can be activated and tested in regular operation without impacting current users.

We plan to complete the design and implementation of the new autonomous configuration scheme, as well as the theoretical analysis, simulation and experimental justification of the scheme, as represented by Mr. Thyagajan's dissertation.

The work on long-range dependency is expected to grow as our understanding of the interesting phenomena grows.

Finally, we plan to complete, present and publish the three papers mentioned above in this report.