EDITOR: David Walden, dave@walden-family.com

DEPARTMENT: ANECDOTES

On the Origin of Kerberos

Jerome H. Saltzer, *Massachusetts Institute of Technology, USA*

Kerberos is a cryptographic authentication and key distribution system developed by MIT Project Athena with the goal that a single login provide access to many different computing services. Kerberos is distributed as a component of most major operating systems, including Microsoft Windows, Apple OS/X and IOS, IBM z/OS, and many versions of Unix. It is also thoroughly documented both in professional papers and in tutorials. However, most descriptions (a typical example is the article in Wikipedia[1]) begin the history of Kerberos by saying it is based on a protocol designed by Roger Needham and Michael Schroeder in 1978. It is not so well known that the Needham & Schroeder protocol is itself based on a 1967 invention by Howard Rosenblum of the United States National Security Agency (NSA) that for many years was classified.

THE KEY DISTRIBUTION PROBLEM

For encrypted communication to be useful, the recipient of a message must be in possession of the decryption key that corresponds to the encryption key used by the sender. In addition, potential adversaries must not be able to get a copy of the decryption key. Thus, cryptographic communication requires advance planning to choose encryption and decryption keys and to distribute those keys to the sender and receiver using a secure communication method. The remainder of this discussion assumes use of symmetric encryption, which means that the encryption and decryption keys are either identical or derivable one from the other, so the secure communication of the keys must assure both confidentiality and integrity.[2]

When in the 1960s the NSA developed a secure telephone system these requirements presented a problem: since anyone with a phone capable of encryption might place a call to any other such phone, it seemed that there had to be an advance arrangement to share a unique encryption key between every potential pair of telephones. For this reason, in the initial implementation, each phone stored a list of encryption keys, one for every other phone it might call or that might call it. Setting up and maintaining these lists presented a scaling problem. Creating a network of a half-dozen secure telephones might be feasible but to deploy thousands for use throughout the Defense Department would be a major challenge, especially since an element of cryptographic security doctrine is that encryption keys should be used only briefly, then changed, to reduce the amount of material encrypted under a single key that could be used for cryptanalysis or that would be exposed if the key were compromised.

For a second generation secure telephone system, Howard Rosenblum in 1967 proposed and patented a key distribution system that he called Bellfield.[3, 4] The idea was that each phone would maintain just one encryption key, unique to that phone, and shared only with a key distribution center (KDC). To place a call from phone A to phone B, the first step would be for phone A to send an encrypted message to the KDC indicating an intent to hold an encrypted conversation with phone B. The KDC would fabricate a new encryption key, known as a "session key" and send two copies of that session key back to A, one encrypted under the key that it shares with A and the second encrypted under the key it shares with B. Phone A then forwards B's encrypted session key as part of the call setup. Once this flurry of messages has been delivered, A and B both possess copies of a session key for this conversation. In addition to allowing A and B to communicate securely without a specific prior arrangement between them, a feature of this technique is that every new telephone conversation could be encrypted with its own unique session key, thus helping satisfy the cryptographic security doctrine to minimize use of any one key.

The early 1970s saw growing interest in both computer security and computer communication networks. In addition to the NSA,[5] several groups working

1058-6180 © 2021 IEEE
Digital Object Identifier 10.1109/MAHC.2021.3051508
Date of current version 5 March 2021.

in industry and academic environments were thinking about how to apply encryption to protect data transmitted between computers. A secure peer-to-peer computer network has a requirement similar to that of a secure telephone network: arranging to share secret encryption keys between a large number of endpoints that may pair unpredictably. Key distribution was an unsolved problem.[6]

While consulting for the Department of Defense in the early 1970s, I learned of Bellfield and realized that this technique, though invented for securing telephones, could be equally applicable to securing computer network communication. But there was a roadblock: The Bellfield technique was classified Confidential and descriptions of it were protected as Sensitive Compartmented Information.

DECLASSIFICATION

Apparently, some folks within NSA also were of the opinion that the technique could be useful in the context of computer communications networks. In 1973, a colleague in the Defense Department alerted me to a newly published paper by NSA staff member Dennis Branstad in the unclassified proceedings of a conference on computer networks held by the American Institute of Aeronautics and Astronautics.[7] Branstad's paper appeared in an unusual venue and most of its content consists of high-level generalities about computer network organization. Although the paper frequently mentions security, identification, authentication, and authorization, the word "encryption" appears just once; the paper instead talks about "keys" for "locking devices" and "unlocking devices" at the communicating endpoints. Then, without mentioning the project name, the origin of the idea, nor its applicability to secure telephone systems, it describes the essence of the Bellfield technique: each potential communicant would share a key with a central agency and the agency would, on request from one party to communicate with another and after checking authorization against an access control list, distribute a session key to the two parties.

Since Branstad had revealed the basic technique of Bellfield, I showed his paper to Roger Needham and Mike Schroeder. When Needham visited the Xerox Palo Alto Research Center in the summer of 1977, he and Schroeder designed a complete computer network protocol for authentication using a KDC. Their protocol combined the technique described by Branstad with a method described by Horst Feistel[8] about using random bit strings, which Needham dubbed "nonces," to assure integrity (that an opponent has not modified a protocol message while it is in transit) and freshness (an opponent has not recorded a previous protocol message and is replaying it).[9] Unaware of the classified origin of Branstad's technique, they distributed their paper in September 1978 as a technical report and three months later published it in the *Communications of the ACM*.[10]

THE PATH TO KERBEROS

The next step was in 1985, when David Clark at MIT supervised two Bachelor of Science theses, by Eric Jaeger and Cliff Neuman, to design a practical key distribution system based on the Needham & Schroeder protocol. Jaeger's thesis added a ticket-granting service to implement Branstad's proposal that the KDC could also perform authorization.[11] Neuman's thesis added discretionary access control and accounting and included implementation of a prototype server.[12]

Upon completing his thesis, Neuman then went on to work with Steve Miller at MIT Project Athena to design and implement the Kerberos authentication system. The protocol they designed includes more efficient protection against attacks and allows relaying of credentials among multiple KDCs operated by independent administrative authorities.[13] They completed the initial implementation of Kerberos, parts of which were based on code from Neuman's thesis, in the fall of 1986 and the authentication system went into production use at Project Athena in January 1987.

In the 34 years since then, cryptography specialists have subjected both the Needham and Schroeder protocol and Kerberos to extensive examination and, to no one's surprise, they have identified several vulnerabilities. In turn, Kerberos has gone through revisions to incorporate techniques to counter these vulnerabilities. But that development is a different story that is deeply technical and well-documented elsewhere.[14]

ACKNOWLEDGMENTS

My thanks go to Mike Schroeder and Cliff Neuman who helped confirm the general narrative and fill in many details, but I claim to have originated any and all mistakes.

REFERENCES/ENDNOTES

1. Wikipedia contributors, "Kerberos," *Wikipedia, The Free Encyclopedia*. Accessed: Aug. 18, 2020. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Kerberos_(protocol)&oldid=977461929

2. An alternative approach uses asymmetric (also known as non-secret or public key) encryption, in which case the requirement is slightly relaxed: The encryption key must be distributed to the sender by a method that assures just integrity. Asymmetric encryption changes many details but the essential problem of key distribution remains and the same invention applies.
3. T. Johnson, "American Cryptology in the Cold War 1945-1949," vol. III, p. 143 (National Security Agency Center for Cryptologic History, 1998, TOP SECRET UMBRA TK, declassified July 26, 2013). Accessed: Dec. 6, 2020. [Online]. Available: https://www.nsa.gov/Portals/70/documents/news-features/declassified-documents/cryptologic-histories/cold_war_iii.pdf
4. H. Rosenblum, "Secure communication system with remote key setting," (U.S. Patent filed in secret Feb. 14, 1969, Jan. 8, 1980. [Online]. Available: https://patents.google.com/patent/US4182933. Original on-line at (search for 4,182,933). Accessed: Aug. 15, 2020. [Online]. Available: http://patft.uspto.gov/netahtml/PTO/srchnum.htm
5. (Author redacted), "NSA history of computer security", iv + 107 pp. (National Security Agency, Feb. 11, 1998, TOP SECRET//COMINT, declassified Sept. 24, 2020). Accessed: Oct. 3, 2020. [Online]. Available: https://cryptome.org/2020/10/nsa-history-computer-security-1998.pdf
6. P. Baran, "On distributed communications: IX. Security, secrecy, and tamper-free considerations," Rand Research Memorandum RM-3765-PR, x + 39 pp. (The Rand Corporation, Aug. 1964). Accessed: Dec. 6, 2020. [Online]. Available:https://www.rand.org/content/dam/rand/pubs/research_memoranda/2006/RM3765.pdf. W. Ware, "Security controls for computer systems: Report of Defense science board task force on computer security," Rand Report R-609, xvi + 68 pp. (The Rand Corporation, Feb. 11, 1970, CONFIDENTIAL, declassified Dec. 31, 1976). Accessed: Dec. 6, 2020. [Online]. Available: https://csrc.nist.gov/csrc/media/publications/conference-paper/1998/10/08/proceedings-of-the-21st-nissc-1998/documents/early-cs-papers/ware70.pdf
7. D. K. Branstad, "Security aspects of computer networks," in *Proc. AIAA Comput. Netw. Syst. Conf.*, Huntsville, Alabama, USA, paper 73-427, p. 5, Apr. 16–18, 1973.
8. H. Feistel, "Cryptographic coding for data-bank privacy," *IBM Research Report RC*, vol. 2827, p. 12 (IBM Corporation, Yorktown Heights, New York, NY, USA, Mar. 18, 1970). Accessed: Oct. 4, 2020. [Online]. Available: https://dominoweb.draco.res.ibm.com/reports/RC2827.pdf
9. Although the Bellfield patent describes parity checking, neither it nor Thomas Johnson's description of Bellfield cited above make mention of protection against modification and replay attacks. Considering that Bellfield was invented within the NSA, it is likely that its implementation included such protections.
10. R. M. Needham and M. D. Schroeder, "Using encryption for authentication in large networks of computers," *Commun. ACM*, vol. 21, 12, pp. 993–999 (Dec. 1978). Accessed: Dec. 6, 2020. [Online]. Available: https://doi.acm.org/10.1145/359657.359659. A prepublication version, Xerox Palo Alto Research Center Tech. Rep. CSL-78-4, 13 pp. (Sept. 1978) has no paywall. Accessed: Dec. 6, 2020. [Online]. Available: www.bitsavers.org/pdf/xerox/parc/techReports/
11. E. Jaeger, "Protocol for trusted third party access control," B.S. thesis, Dept. Elect. Eng. Comp. Sci., MIT, Feb. 1985.
12. B. C. Neuman, "Sentry: A discretionary access control server," B.S. Thesis, Dept. Elect. Eng. Comp. Sci, MIT, Cambridge, MA, USA, May 1985.
13. S. P. Miller, B. C. Neuman, J. I. Schiller, and J. H. Saltzer, "Section E.2.1: Kerberos authentication and authorization system," Project Athena Technical Plan, MIT Project Athena, Cambridge, MA, USA, 36 pp. (Oct. 1988). Accessed: Dec. 6, 2020. [Online]. Available: http://web.mit.edu/Saltzer/www/publications/athenaplan/e.2.1.pdf
14. See, for example, the following paper and the 40 papers it cites: T. Yu, S. Hartman, and K. Raeburn, "The perils of unauthenticated encryption: Kerberos version 4," in *Proc. Network Distrib. Syst. Secur. Symp.*, San Diego, CA, USA, (The Internet Society, Feb. 5, 2004). Accessed: Oct. 6, 2020. [Online]. Available: https://www.ndss-symposium.org/ndss2004/perils-unauthenticated-encryption-kerberos-version-4/

JEROME H. SALTZER is a Professor Emeritus of Computer Science, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science. At the time of the events described here, he was the Head of the Computer Systems Research Group, MIT Laboratory for Computer Science. Contact him at Saltzer@mit.edu.