Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2017

S. Goldberg
Boston University
H. Stenn
Network Time Foundation

July 8, 2016

Network Time Protocol Not You REFID
draft-stenn-ntp-not-you-refid-00

Abstract

NTP has been widely used through several revisions, with the latest being RFC 5905 [RFC5905]. A core component of the protocol and the algoritms is the Reference ID, or REFID, which is used to identify the source of time used for synchronization (aka the "system peer"). Traditionally, when the source of time was another system, the REFID was the IPv4 address of that other system. The purpose of the REFID is to prevent a one-degree "timing loop": where if A has several timing sources that include B, if B decides to get its time from A, then A should not then decide to get its time from B. The REFID is therefore a vital core-component of the base NTP packet. If a system's REFID is the IPv4 address of its time source, then with a simple query a remote attacker can learn the target's REFID. The remote attacker can then try to use that information to send spoofed NTP packets to the target or the target's time source, attempting to cause a disruption in time service [NDSS16]. Since the core purpose of the REFID is to prevent a one-degree timing loop, this proposal is a backward-compatible way to limit the amount of information that is leaked in the REFID. Specifically, it allows the prevention of one-degree timing loops by allowing a system A to reveal to a querying system B that B is not A's time source, but without revealing the actual time source to which A is synchronized.

This Internet-Draft will expire on January 9, 2017.

Table of Contents

1. Introduction

NTP has been widely used through several revisions, with the latest being RFC 5905 [RFC5905]. A core component of the protocol and the algoritms is the Reference ID, or REFID, which is used to identify the source of time used for synchronization. Traditionally, when the source of time was another system, the REFID was the IPv4 address of that other system. If the source of time was using IPv6 for its connection, then a 4 octet digest value of the IPv6 address was used as the REFID. The purpose of the REFID is to prevent a one-degree timing loop, where if A has several timing sources that include B, if B decides to get its time from A, then A should not then decide to get its time from B.

Recently it was observed in [NDSS16] that a remote attacker can query a target system to learn its time source from the REFID included in target's response packet. The remote attacker can then use this information to send spoofed packets to the target or its time source,

in an attempt to disrupt time service. The REFID thus unnecessarily leaks information about a target's time server to remote attackers. The best way to mitigate this vulnerability is to decouple the IP address of the time source from the REFID. To do this, a system can use an otherwise-impossible value for its REFID, called the "not-you" value, when it believes that a querying system is not its time source.

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. The REFID

The interpretation of a REFID is based on the stratum, as documented in RFC 5905 [RFC5905], section 7.3, "Packet Header Variables". The core reason for the REFID in the NTP Protocol is to prevent a degree-one timing loop, where server B decides to follow A as its time source, and A then decides to follow B as its time source.

At Stratum 2+, which will be the case if two servers A and B are exchanging timing information, then if server B follows A as its time source, A's address will be B's REFID. When A uses IPv4, the default REFID is A's IPv4 address. When A uses IPv6, the default REFID is a four-octet digest of A's IPv6 address. Now, if A queries B for its time, then A will learn that B is using A as its time source by observing A's address in the REFID field of the response packet sent by B. Thus, A will not select B as a potential time source, since this would cause a timing loop.

However, this mechanism also allows a third-party C to learn that A is the time source that is being used by B. When A is using IPv4, C can learn this by querying B for its time, and observing that the REFID in B's response is the IPv4 address of A. Meanwhile, when A is using IPv6, then C can again query B for its time, and then can use an offline dictionary attack to determine the IPv6 address that corresponds to the digest value in the response sent by B. C could construct the necessary dictionary by compiling a list of publically accessible IPv6 servers. Remote attackers can use this technique to identify the time sources used by a target, and then send spoofed packets to the target or its time source in order to disrupt time service as was done e.g., in [NDSS16] or [CVE-2015-8138].

3. The Not-You REFID

This proposal allows the one-degree loop detection to work while keeping potentially abusable information from being disclosed to uninterested parties. It does this by returning the normal REFID to queries that come from an address that the current system believes is its time source (aka its "system peer"), and otherwise returning a special IP address that is interpreted to mean "not you". The "not you" IP address is 127.127.127.127 when the query is made from an IPv4 address, or when the query is made from an IPv6 address whose four-octect hash does not equal 127.127.127.127. The "not you" IP address is 127.127.127.128 when the query is made from an address whose four-octect hash equals to 127.127.127.127.

Note that this mechanism is transparent to the party that sends timing queries. A querying system that uses IPv4 continues to check that its IPv4 address does not appear in the REFID before deciding whether to take time from the current system. A querying system that uses IPv6 continues to check that the four-octet hash of its IPv6 address does not appear in the REFID before deciding whether to take time from the current system.

This proposal will hide the current system's system peer from querying systems that the current system believes are not the current system's system peer. Note well, however, that the current system will return the "not you" value to a query from its system peer if the system peer sends its query from an unexpected IP address.

4. Security Considerations

Many systems running NTP are configured to return responses to timing queries by default. These responses contain a REFID field, which generally reveals the address of the system's time source. This behavior can be exploited by remote attackers who wish to first learn the address of a target's time source, and then attack the target and/or its time source. As such, this proposal is designed to harden NTP against these attacks by limiting the amount of information leaked in the REFID field.

5. Acknowledgements

The authors wish to acknowledge useful discussions with Aanchal Malhotra and Matthew Van Gundy.

6. References

6.1. 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>.

6.2. Informative References

[CVE-2015-8138] Van Gundy, M. and J. Gardner, "Network Time Protocol Origin Timestamp Check Impersonation Vulnerability (CVE-2015-8138)", in TALOS VULNERABILITY REPORT (TALOS-2016-0077), 2016.

[NDSS16] Malhotra, A., Cohen, I., Brakke, E., and S. Goldberg, "Attacking the Network Time Protocol", in ISOC Network and Distributed System Security Symposium 2016 (NDSS'16), 2016.

Authors' Addresses

Sharon Goldberg
Boston University
111 Cummington St
Boston, MA 02215
US

Email: goldbe@cs.bu.edu

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

Email: stenn@nwtime.org