INTERNET DRAFT
draft-ohta-practically-secure-dns-00.txt
Intended status: Standard Track
Expires: April 23, 2012

M. Ohta
Tokyo Institute of Technology

October 24, 2011

Practically Secure DNS

Status of this Memo

Copyright Notice

Abstract

Plain DNS without PKI is secure, if a chain of query/response communications between a client and an authoritative server relayed by zero or more intermediate resolvers and the authoritative server and all the resolvers are secure.

However, because of short (16bit) message ID, the communications composing the chain are not very secure without, or even with (port exhaustion attack is possible), source port randomization.

Still, plain DNS can be made practically secure, if the client makes two queries with independent message IDs to an address of a server (a resolver or a name server) and confirm that two replies are identical.

1. Security Considerations

This memo proposes rather an implementation guideline than protocol modifications to DNS [DNS] to make DNS without PKI practically secure by letting a DNS client make two (or more, if necessary) queries, which are identical save cryptographically randomized message IDs, to a server (a resolver or a name server) over UDP and let the client check whether the replies are identical or not, ignoring inessential differences of RR ordering, character cases of labels and compression of labels.

Queries should have identical source and destination addresses and randomized source port numbers. The second query may immediately follow the first one without a randomized interval.

While multiple queries increases the number of packets exchanged a little more than twice, the number of packets and associated computation load are expected to be considerably less than queries/replies over TCP (according to [TTCP], 10 packets are exchanged) and much less than those with [SDNS].

If there is no reply after reasonable time out, it is a temporal failure.

If there is only one reply after the reasonable time out, another query should be sent, failure to get the second reply after the reasonable time out is a temporal failure.

If two replies are received but not identical, it should be because:

1) versions of some zones related to the queries changes between the queries; or

2) replies are constructed from different cache content; or

3) the server is actively randomizing answers for load balancing and other purposes; or

4) the client is under a message ID guessing attack.

Even though two queries are made with a short interval, 1) and 2) can still occur, because zone version and cache content may change during the interval.

In any case, additional one or two queries, with different message IDs, should be generated.

For cases 1), 2), 3) (only if randomization is based on client address) and 4), one or two more queries should be enough to have two identical replies, which should be practically secure. If no identical replies are available, it should be case 3) with truly random replies.

During waiting replies, case 4) can be identified by replies with unmatched message IDs, which should initiate serious security alerts.

Otherwise, it should be case 3) and, unless case 4) is identified to be simultaneously occurring, one of the replies should be practically secure, though it is less secure. As DNS assumes "the system should be able to deal with subsets that change more rapidly (on the order of seconds or minutes)" [DNS], it is a natural consequence of randomizing servers to break the fundamental principle of DNS. If case 3) and 4) is simultaneously occurring, it's a temporal failure.

A resolver used by a practically secure DNS client should also behave as a client of practically secure DNS to make DNS practically secure end to end. However, then, the resolver must merge identical queries from its clients to prevent exponential growth of DNS traffic. If the resolver consists of internal multiple servers, identical queries should be forwarded to a single internal server based on source addresses of the queries, to prevent the exponential growth.

A resolver unaware of practically secure DNS may generate multiple queries as a client, if corresponding multiple queries come from its clients.

2. IANA Considerations

This document has no actions for IANA.

Normative References

[DNS] STD0013.

Informative References

[TTCP] R. Braden, "Extending TCP for Transactions -- Concepts", . RFC1379, November 1992.

[SDNS] R. Arends, R. Austein, M. Larson, D. Massey and S. Rose, "DNS Security Introduction and Requirements". RFC4033 , March 2005.

Author's Address

Masataka Ohta
Graduate School of Information Science and Engineering
Tokyo Institute of Technology
2-12-1, O-okayama, Meguro-ku
Tokyo 152-8552, JAPAN

Phone: +81-3-5734-3299
Fax: +81-3-5734-3299
EMail: mohta@necom830.hpcl.titech.ac.jp