IETF dhc Working Group T. Jinmei
Internet-Draft Toshiba
Expires: December 28, 2006 June 26, 2006

Clarifications on DHCPv6 Authentication
draft-ietf-dhc-dhcpv6-clarify-auth-01.txt

Abstract

This document describes issues about the authentication mechanism of the Dynamic Host Configuration Protocol for IP version 6 that were identified from implementation experiences. It also tries to propose resolutions to some of the issues.

1. Introduction

Several questions arose on the authentication mechanism of DHCPv6

[RFC3315] from implementation experiences, particularly on its delayed authentication protocol. Some of the questions may require a change or addition to the current protocol, and one of them may even cause discussions on a security threat.

This document describes the issues based on the questions and proposes resolutions, hoping the resolutions will be merged in if valid and accepted, to the next version of the base specification.

2. Usage with Information-Request

According to [RFC3315], it seems possible to use the authentication mechanism for Information-request and Reply exchanges. The RFC says in Section 21.4.4.4 as follows:

> If the server has selected a key for the client in a previous message exchange (see section 21.4.5.1), the client MUST use the same key to generate the authentication information throughout the session.

However, this description is not really clear. Section 21.4.5.1, which is referred from the above part, actually describes the case of Solicit and Advertise exchange:

> 21.4.5.1. Receiving Solicit Messages and Sending Advertise Messages
>
> The server selects a key for the client and includes authentication information in the Advertise message returned to the client as specified in section 21.4. [...]

It does not necessarily mean contradiction because the client and the server may have exchanged Solicit and Advertise with authentication before starting the Information-request and Reply exchange. But it then restricts the usage scenario of the authentication mechanism for Information-request and Reply exchanges. In particular, this assumption prohibits the use of the mechanism with the "stateless" service using DHCPv6 [RFC3736]. Whereas the specification allows an implementation that only supports the stateless service and does not support Solicit and Advertise messages, the authentication mechanism depends on Solicit and Advertise exchanges.

This fact can (partly) invalidate a security consideration in [RFC3736]:

> Authenticated DHCP, as described in sections 21 and 22.11 of the DHCP specification [1], can be used to avoid attacks mounted through the stateless DHCP service.

(where [1] refers to [RFC3315].) In fact, as was just shown above, authenticated DHCP cannot be used unless the implementations also support Solicit and Advertise messages (or the entire [RFC3315] in general).

It should also be noted that [RFC3315] does not define what the server should do when it receives an Information-request message containing an authentication option; Section 21.4.5.2 excludes the Information-request message.

2.1. Suggested Resolution

Considering the fact that [RFC3736] allows implementations that only support the subset of the full specification [RFC3315], it should make sense to define the authentication usage for Information-request and Reply exchanges separately.

One significant difference between Information-request and other "stateful" cases is that there is no explicit notion of "session" in the former. In some cases, however, the same client and server may exchange Information-request and Reply multiple times, where the entire exchanges can be regarded as a "session". For example, the client may want to get different configuration information in multiple exchanges. Also, if the client and the server use the Information Refresh Time Option [RFC4242], they will restart exchanges when the refresh time expires.

A naive implementation of keeping a "session" in the server will decrease an advantage of the [RFC3736] usage that the server can run in a stateless fashion without any client-specific state. Specifically, the server may have to maintain the following three types of state per client:

- the authentication key shared between the server and the client
- replay detection information for messages to the client (such as the replay detection value sent most recently)
- replay detection information for messages from the client (such as the replay detection value received most recently)

It is possible for the server to have different keys for different clients without having per-client information by the method described in Appendix A of [RFC3118]. In this approach the server only maintains a single master key and creates the key for a particular

client by computing some one-way hash digest based on the master key and the client's DUID. Since an Information-request message to be authenticated must have a Client Identifier option as specified in Section 18.1.5 of [RFC3315], this approach should work well.

The server can also avoid a replay attack using an old message sent from the server without maintaining per client state for replay detection information about messages sent to the client if the server has a source of replay protection value that monotonically increases. For example, system timestamp can be used for this purpose.

Without keeping the replay detection information for messages from the client, the server may be vulnerable to a replay attack from a malicious client. This should be relatively a minor issue because a "stateless" server usually provides the same information for all clients without consuming any of its resource. When the malicious client can get an old valid message, e.g., by snooping the traffic, it would also be able to get and use the response; it does not have to mount the replay attack.

There is still a subtle case to be considered. If the server uses a monotonically increasing counter for stateless replay detection information to clients and the server does not maintain per client replay detection information from the client, a malicious client can reuse a valid Information-request message to get a reusable valid Response message. The malicious client will then be able to mount a replay attack on the client later.

The proposed revision of Section 21.4.4.4 is therefore as follows:

21.4.4.4. Sending Information-request Messages

When the client sends an Information-request message and wishes to use authentication, it includes an Authentication option with the desired protocol, algorithm and RDM as described in section 21.4. The client does not include any replay detection or authentication information in the Authentication option.

If the client authenticated exchanges of Information-request and Reply in the past, the client MAY reuse the same key used in the previous exchanges to generate the authentication information. In this case, the client generates authentication information for the Information-request message as described in section 21.4.

Note that the keys used for these exchanges are separately managed from the keys used for the other exchanges beginning with the Solicit message when the two types of exchanges run concurrently, while the two keys may happen to be the same. For example, replay

detection should be performed separately, and validation failure for one type of exchanges does not affect the other.

Section 21.4.4.5 will also need to be revised. However, since this section has a separate issue per se as will be discussed in Section 5, further details on this are not discussed here.

The server side behavior needs to be described, too. Along with the above change to Section 21.4.4.4, a new proposed subsection of Section 21.4.5 should be added as follows:

21.4.5.x. Receiving Information-request Messages and Sending Reply Messages

If the Information-request message includes an authentication option without authentication information, the server selects a key for the client and includes authentication information in the Reply message returned to the client as specified in section 21.4. If the key is selected from pre-configured information for the client maintained in the server, the server MUST record the identifier of the key selected for the client so that it can validate further Information-request messages from the client if

the client reuses the same key for the future exchanges.

The server MAY alternatively use a stateless method such as the one described in Appendix A of [RFC3118]. Then the server MUST consistently use the same key for the client to validate further Information-request messages from the client.

When the server uses such a stateless method for key utilization, it may also seek to avoid having per client state for replay detection. For outbound messages, it is easy when the server has a source of monotonically increasing values such as system timestamp; however, it is difficult if not impossible to refuse inbound replay messages without per client state. The server MAY skip the replay detection for inbound Information-request messages if the benefit of being stateless outweighs the risk of replay attacks for inbound messages. Care should be taken when this relaxation applies because if the server also uses a stateless method for outbound messages, a malicious client may be able to get a valid reusable response by reusing an old, legitimate Information-request message.

If the Information-request message includes an authentication option with authentication information, the server uses the key identified in the message and validates the message as specified in section 21.4.2. If the message fails to pass the validation test, or the key identified by the authentication information of

the message is not identical to the key that the server used in the previous exchange (when it has recorded the key), the server MUST discard the message and MAY choose to log the validation failure.

If the message passes the validation test, the server responds with a Reply message as described in section 18.2.5. The server MUST include authentication information generated using the key just selected or identified in the received message, as specified in section 21.4.

Finally, if the server previously recorded the key but receives an Information-request message without including an authentication option, the server MUST accept the message and respond with a Reply message without including authentication information. The

server SHOULD then remove the recorded information.

Note that the keys used for these exchanges are separately managed from the keys used for the other exchanges beginning with the Solicit message when the two types of exchanges run concurrently (See Section 21.4.4.4).

3. What If Replay Is Detected

It is not clear what the receiver should do when an attempt of replay attack is detected from either Section 21.3 or Section 21.4.2 of [RFC3315].

3.1. Suggested Resolution

It should be natural to discard a DHCP message containing an authentication option whose replay detection field indicates a replay attack.

Instead of concentrating on this particular case, we propose to revise the entire second paragraph of Section 21.4.2 as follows:

To validate an incoming message, the receiver first checks that the value in the replay detection field is acceptable according to the replay detection method specified by the RDM field. If no replay is detected, then the receiver computes the MAC as described in [8]. The entire DHCP message (setting the MAC field of the authentication option to 0) is used as input to the HMAC-MD5 computation function. If the MAC computed by the receiver matches the MAC contained in the authentication option, the message regarded as valid. If the above procedure fails at any stage, the receiver MUST discard the DHCP message.

4. Inconsistent Behavior for Unauthenticated Messages

[RFC3315] says in Section 21.4.2 (Message Validation) as follows:

If the MAC computed by the receiver does not match the MAC contained in the authentication option, the receiver MUST discard the DHCP message.

On the other hand, Section 21.4.4.2 allows the client to respond to an Advertise even if it fails to authenticate the message:

> Client behavior, if no Advertise messages include authentication information or pass the validation test, is controlled by local policy on the client. According to client policy, the client MAY choose to respond to an Advertise message that has not been authenticated.

This seems to say, for example, that the client MAY accept an Advertise message based on its local policy, even if the MAC computed by the receiver does not match the MAC contained in the authentication option. Apparently this contradicts the requirement in Section 21.4.2.

4.1. Suggested Resolution

There seems to be no valid reason for accepting an Advertise message if it fails validation. On the other hand, it may make sense in some cases that the client accepts messages that do not include an authentication option or even messages with an authentication option which specifies a key the client does not know.

As a related terminology issue, [RFC3315] uses the phrase of "unauthenticated message(s)" in Sections 21.4.4.2 and 21.4.4.5 without formally defining the term. Based on the above discussion, the most appropriate definition of this term is the acceptable types of messages. Specifically, a message that fails to pass the validation test should not be regarded as an "unauthenticated" message.

The suggested change to Section 21.4.4.2 is thus as follows.

> The client validates any Advertise messages containing an Authentication option specifying the delayed authentication protocol using the validation test described in section 21.4.2.
>
> Client behavior, if all Advertise messages are "unauthenticated", is controlled by local policy on the client, where an unauthenticated message means a message that does not include an

> authentication option or a message with an authentication option

> which specifies a key the client does not know. According to client policy, the client MAY choose to respond to an unauthenticated Advertise message.
>
> [...]
>
> A client MUST be configurable to discard unauthenticated messages, and SHOULD be configured by default to discard unauthenticated messages if the client has been configured with an authentication key or other authentication information. A client MAY choose to differentiate between unauthenticated Advertise messages with no authentication information and unauthenticated Advertise messages that specifies a key the client does not know; for example, a client might accept the former and discard the latter. If a client does accept an unauthenticated message, the client SHOULD inform any local users and SHOULD log the event.

The second paragraph of Section 21.4.4.5 also needs a change. But, again, we will discuss this case in Section 5.

5. Possibility of DoS Attack

Section 21.4.4.5 of the RFC says as follows:

> If the Reply fails to pass the validation test, the client MUST restart the DHCP configuration process by sending a Solicit message.

The purpose of this specification is probably to avoid a deadlock scenario when the server suddenly reboots forgetting the authentication key and/or the replay detection counter.

However, this behavior can easily cause denial of service (DoS) attacks; the attacker can simply send an invalid Reply message at some valid timing and can invalidate configuration information of the client or can prevent the client from configuring itself.

As a side issue, this section seems to not consider Information-request and Reply exchanges.

5.1. Discussion on Resolution

Even if a Reply message does not pass the validation tests, it is probably reasonable to wait a certain period for an authenticated one. Additionally, if the Reply message is a response to Release, the client will not have to restart the configuration process with

Solicit. It can simply quit the session after the waiting period. The appropriate waiting period would be the first timeout for the message, since in the intended scenario described above the legitimate (but without knowing a valid key) server should be working and respond within the timeout period.

Reply messages to Information-request will need a separate consideration. According to the resolution (to a different issue) suggested in Section 2.1, the client may or may not reuse the key for previous exchanges. If the client does not reuse the key, or if this is the first time the client sends an Information-request message, there should be no other behavior than simply discarding the message and waiting for a valid response (usual timeout and resend will apply). Otherwise, the appropriate behavior would be similar to the case described in the previous paragraph. That is, the client should wait for a while and start new exchanges without including authentication information with the reused key.

5.2. Suggested Resolution

The suggested change to Section 21.4.4.5 based on the above analysis is as follows. It includes resolutions to the issues discussed in Section 2.1 and Section 4.1.

> For Reply to a message other than Information-request, if the client authenticated the Advertise it accepted, the client MUST validate the associated Reply message from the server. The client MUST discard the Reply if the message fails to pass the validation test and MAY log the validation failure. Then the client MUST wait until the corresponding timeout expires as specified in Section 14 as if it did not receive any Reply. If the client cannot receive a valid Reply within the first timeout period, the client MUST restart the DHCP configuration process by sending a Solicit message or the client MUST simply quit the configuration process if the Reply should be a response to a Release message.
>
> If the client accepted an unauthenticated Advertise message that did not include authentication information or did not pass the validation test, the client MAY accept an unauthenticated Reply message from the server.
>
> If the client sent an Information-request message including an authentication option, the client MUST validate the associated Reply message from the server. The client MUST discard the Reply if the message fails to pass the validation test and MAY log the validation failure. If the client reuses the key used in previous

exchanges and authenticated the Information-request message with the key, the client MUST wait until the corresponding timeout

expires. If the client cannot receive a valid Reply within the first timeout period, the client MUST restart the configuration process by restarting exchanges of Information-request and Reply without reusing the previous key.

The client MAY choose to accept an unauthenticated Reply message to an Information request message. All the discussions and behaviors described in Section 21.4.4.2 should simply apply to this case.

6. Lack of Authentication from Client

It is not clear what the server should do when the client does not include an authentication option while the server has previously sent authentication information in the same session.

A proposal for the Information-request message was provided in Section 2.1. For messages other than Information-request, the lack of an authentication information would mean the Advertise message sent from the server was "unauthenticated" to the client and the client chose to accept that message. Thus, the server should basically accept that message, while it may still want to reject the message if the server uses the authentication information to authenticate the client.

The proposed change to Section 21.4.5.2 based on the above discussion is to add the following paragraph at the end of the section.

If the message does not include an authentication option while the server included an authentication option in previous messages of the session, the server SHOULD accept the message. In this case, the server MAY skip including an authentication option in the Reply message. The server MAY still choose to discard the received message in this case based on its local policy.

7. Key Consistency

[RFC3315] requests in Section 21.4.4.3 that the client use the same key used by the server to generate the authentication information. However, it is not clear from the RFC what the server should do if the client breaks this rule. It says in Section 21.4.5.2 that

> If the message [...] or the server does not know the key identified by the 'key ID' field, the server MUST discard the message and MAY choose to log the validation failure.

It is not clear whether "does not know the key" means a different key from the one the server specified in the Advertise message. If this is the intent, this sentence should be clarified as follows:

> If the message [...] or the key identified by the authentication information of the message is not identical to the key that the server has been using in the session, the server MUST discard the message and MAY choose to log the validation failure.

8. Security Considerations

This document specifically talks about security issues for DHCPv6. It also points out a possibility of DoS attacks, and proposes a resolution to prevent the attacks.

9. Acknowledgements

Christian Strauf, Stig Venaas and Bernie Volz reviewed a preliminary version of this document, and provided specific suggestions and further clarifications.

10. IANA Considerations

This document has no actions for IANA.

11. References

11.1. Normative References

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.

[RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.

11.2. Informative References

[RFC4242] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, November 2005.

Appendix A. CHANGE HISTORY

Changes since draft-jinmei-dhc-dhcpv6-clarify-auth-00.txt are:

- Loosened the description of Information-request and Reply exchanges so that the server can skip maintaining per-client information.
- Changed the lifetime option to Information Refresh Time Option, according to the change of the document.
- Clarified the definition of "unauthenticated" messages so that those would not include messages that failed the validation, and revise the specification using the term. Section 4 of the previous version was removed accordingly.
- Provided specific suggestion to solve denial of service attacks.
- Changed the draft name to an IETF dhc working group document.

Changes since draft-ietf-dhc-dhcpv6-clarify-auth-00.txt are:

- Updated the considerations and recommendation on the server behavior for Information-request and Reply exchanges based on stateless key utilization described in [RFC3118].

Author's Address

Tatuya Jinmei
Corporate Research & Development Center, Toshiba Corporation
1 Komukai Toshiba-cho, Saiwai-ku
Kawasaki-shi, Kanagawa 212-8582
Japan

Phone: +81 44-549-2230
Email: jinmei@isl.rdc.toshiba.co.jp

Intellectual Property Statement

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.