

Network Working Group
Internet Draft
Intended status: Proposed Standard
Expires: August 28, 2013

B. Liu
Huawei Technologies
R. Bonica
Juniper Networks
February 25, 2013

DHCPv6/SLAAC Address Configuration Interaction Problem Statement
draft-liu-bonica-dhcpv6-slaac-problem-01.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 28, 2013.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

This document analyzes the host behavior of DHCPv6/SLAAC interaction issue. It reviews the standard definition of the host behaviors and

provides the test results of current mainstream implementations. Some potential operational gaps of the interaction are also described.

Table of Contents

1.	Introduction	3
2.	Host Behavior of DHCPv6/SLAAC Interaction	3
2.1.	Relevant RA Flags Defined in Standards	4
2.1.1.	A (Autonomous) Flag	4
2.1.2.	M (Managed) Flag	4
2.1.3.	O (Otherconfig) Flag	4
2.2.	Behaviors of Current Implementations	5
2.2.1.	A flag	5
2.2.2.	M flag	5
2.2.3.	O flag	6
3.	Possible Operational Gaps of DHCPv6/SLAAC Interaction	6
3.1.	Renumbering	6
3.2.	Cold Start Problems	7
3.3.	Strong Management	7
4.	Conclusions	7
5.	Security Considerations	7
6.	IANA Considerations	7
7.	References	8
7.1.	Normative References	8
7.2.	Informative References	8
8.	Acknowledgments	8
Appendix A.	Test Details of Host Behaviors	10
A.1	Host Initialing Behavior	10
A.2	Host SLAAC/DHCPv6 Switching Behavior	11
A.3	Host Stateful/Stateless DHCPv6 Behavior	12
	Authors' Addresses	13

1. Introduction

In IPv6, both of the DHCPv6 [[RFC3315](#)] and Neighbor Discovery [[RFC4861](#)] protocols can provide automatic IP address configuration for the hosts. They are known as stateful address auto-configuration and SLAAC (stateless address auto-configuration)[[RFC4862](#)], and are suitable for different scenarios respectively. Sometimes the two address configuration modes may be both available in one network.

In ND protocol, there is a M (ManagedFlag) flag defined in RA message, indicating the hosts there is DHCPv6 service in the network if the flag is set. And there is an O "OtherConfigFlag", if set, indicating configure information other than addresses (e.g. DNS, Route .etc) is available through DHCPv6 configuration. Moreover, there's another A (Autonomous) flag defined in ND, which indicating the hosts to do SLAAC, may also influent the behavior of hosts.

So with the A/M/O flags, the two separated address configuration modes are somehow correlated. But for some reason, the ND protocol didn't define the flags as prescriptive but only advisory. This ambiguous definition may vary the behavior of hosts when interpreting the flags. In [section 2](#), we provided a brief test result to identify different host operating systems have taken different approaches. This would add additional complexity for both the hosts and the network management.

This draft reviews the standard definition of the above mentioned flags, and provides a test result of several major desktop operating systems' behavior. And then identifies potential requirement/gaps of DHCPv6/SLAAC interaction.

2. Host Behavior of DHCPv6/SLAAC Interaction

In this section, we analyzed A/M/O flags definition, and provide the test result of host behavior of interpreting these flags in mainstream operating systems implementations.

Please note that, A flag has no direct relationship with DHCPv6, but it is somewhat correlated with M/O flags.

2.1. Relevant RA Flags Defined in Standards

2.1.1. A (Autonomous) Flag

In ND Prefix Information Option, the autonomous address-configuration flag (A flag). When set indicates that this prefix can be used for stateless address configuration as specified in SLAAC.

For the host behavior, there is an explicit rule in the SLAAC specification [[RFC4862](#)]: "If the Autonomous flag is not set, silently ignore the Prefix Information option."

But when A flag is set, the SLAAC protocol didn't provide a prescriptive definition.

2.1.2. M (Managed) Flag

In earlier SLAAC specification [[RFC2462](#)], the host behavior of interpreting M flag is as below:

"On receipt of a valid Router Advertisement, a host copies the value of the advertisement's M bit into ManagedFlag. If the value of ManagedFlag changes from FALSE to TRUE, and the host is not already running the stateful address autoconfiguration protocol, the host should invoke the stateful address auto-configuration protocol, requesting both address information and other information. If the value of the ManagedFlag changes from TRUE to FALSE, the host should continue running the stateful address auto-configuration, i.e., the change in the value of the ManagedFlag has no effect. If the value of the flag stays unchanged, no special action takes place. In particular, a host MUST NOT reinvoke stateful address configuration if it is already participating in the stateful protocol as a result of an earlier advertisement."

But in the updated SLAAC specification [[RFC4862](#)], the relative description was removed, the reason was "considering the maturity of implementations and operational experiences. ManagedFlag and OtherConfigFlag were removed accordingly. (Note that this change does not mean the use of these flags is deprecated.)"

2.1.3. O (Otherconfig) Flag

As mentioned above, the situation of O flag is similar with M. In earlier SLAAC [[RFC2462](#)], the host behavior is clear:

"If the value of OtherConfigFlag changes from FALSE to TRUE, the host should invoke the stateful autoconfiguration protocol, requesting

information (excluding addresses if ManagedFlag is set to FALSE). If the value of the OtherConfigFlag changes from TRUE to FALSE, the host should continue running the stateful address autoconfiguration protocol, i.e., the change in the value of OtherConfigFlag has no effect. If the value of the flag stays unchanged, no special action takes place. In particular, a host MUST NOT reinvoke stateful configuration if it is already participating in the stateful protocol as a result of an earlier advertisement."

And there's another description of the relationship of M and O flags in [[RFC2462](#)]:

"In addition, when the value of the ManagedFlag is TRUE, the value of OtherConfigFlag is implicitly TRUE as well. It is not a valid configuration for a host to use stateful address autoconfiguration to request addresses only, without also accepting other configuration information."

[2.2. Behaviors of Current Implementations](#)

We did tests of current 3 mainstream desktop operating systems on the behaviors; please refer to the appendix for details. This section illustrates the important results of the tests.

[2.2.1. A flag](#)

A flag is a switch to control whether to do SLAAC, and it is independent with M/O flags, in another word, A is independent with DHCPv6.

At the non-SLAAC-config state (either non-configured or DHCPv6-configured only), the 3 OSes acted the same with A flag, if A set, they all configured SLAAC, it is obvious and reasonable. But when SLAAC-configured, and A changed from 1 to 0, the behaviors varied, some deprecated SLAAC while some ignored the RA messages.

[2.2.2. M flag](#)

M is a key flag to interact ND/DHCPv6, but the host behaviors on M flag were quite different.

In our test, one OS treats the flag as instruction, it even released DHCPv6 session when M=0. But the other two just treat the flag as advisory, when SLAAC was done, it won't care about M=1, and M=0 won't cause operation for the already configured DHCPv6 addresses. Moreover, the two OSes even would not initiate DHCPv6 session until they

receives RA messages with M=1, this behavior has an implication that DHCPv6 somehow depends on ND.

Please refer to [I-D.liu-6renum-dhcpv6-slaac-switching] for more details.

2.2.3. 0 flag

In our tests, when M flag is set, the O flag is implicitly set as well; in another word, the hosts would not initial stateful DHCPv6 and stateless DHCPv6 respectively. This is a reasonable behavior.

But the O flag is not independent from A flag in some OSes. In our test, there are two OSes won't initiate stateless DHCPv6 when A flag is not set, that is to say, it is not applicable to have a "stateless DHCPv6 only" configuration state for some operating systems; it is also not applicable for these two OSes to switch between stateful DHCPv6 and stateless DHCPv6 (according to O flag changing from 0 to 1 or verse vice).

3. Possible Operational Gaps of DHCPv6/SLAAC Interaction

According to the abovementioned tests, there are possible operational issues as the following.

3.1. Renumbering

During IPv6 renumbering, the SLAAC-configured hosts can reconfigure IP addresses by receiving ND Router Advertisement (RA) messages containing new prefix information. The DHCPv6-configured hosts can reconfigure addresses by initialing RENEW sessions when the current addresses' lease time is expired or receiving the reconfiguration messages initialed by the DHCPv6 servers.

The above mechanisms have an implicit assumption that SLAAC-configured hosts will remain SLAAC while DHCPv6-managed hosts will remain DHCPv6-managed. But in some situations, SLAAC-configured hosts may need to switch to DHCPv6-managed, or verse vice. In [I-D.ietf-6renum-enterprise], it described several renumbering scenarios in enterprise network for this requirement; for example, the network may split, merge, relocate or reorganize. But due to current implementations, this requirement is not applicable and has been identified as a gap in [I-D.ietf-6renum-gap-analysis].

3.2. Cold Start Problems

If all nodes, or many nodes, restart at the same time after a power cut, the results might not be consistent.

3.3. Strong Management

Since the host behavior of address configuration is somehow uncontrolled by the network side, it might cause gaps to the networks that need strong management (for example, the enterprise networks and the ISP CPE networks). Examples are:

- the network wants the hosts to do DHCPv6-only configuration, it is not applicable for some operating systems due to current implementation unless manually configure the hosts to DHCPv6-only model
- the hosts have been SLAAC-configured, then the network needs the hosts to do DHCPv6 simultaneously (e.g. for multihoming)
- the network wants the hosts to do stateless DHCPv6-only; for example, the hosts are configured with self-generated addresses (e.g. ULA), and they also need to contact the DHCPv6 server for information configuration

4. Conclusions

- The host behavior of SLAAC/DHCPv6 interaction is ambiguous in standard.
- The implementations have been varied on this issue. In [\[RFC4862\]](#) it is said "Removed the text regarding the M and O flags, considering the maturity of implementations and operational experiences." The description seems not true anymore.
- It is foreseeable that the un-uniformed host behavior can cause operational gaps, e.g. in renumbering and strong management.

5. Security Considerations

No more security considerations than the Neighbor Discovery protocol [\[RFC4861\]](#).

6. IANA Considerations

None.

7. References

7.1. Normative References

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), September 2007.

7.2. Informative References

- [RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", [RFC 2462](#), December 1998.
- [RFC3315] R. Droms, 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.
- [RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", [RFC 5887](#), May 2010.
- [I-D.ietf-6renum-gap-analysis]
Liu, B., and Jiang, S., "IPv6 Site Renumbering Gap Analysis", Working in Progress, March 2012
- [I-D.ietf-6renum-enterprise]
Jiang, S., and B. Liu, "IPv6 Enterprise Network Renumbering Scenarios and Guidelines ", Working in Progress, March 2012.

8. Acknowledgments

The test was done by our research partner BNRC-BUPT (Broad Network Research Centre in Beijing University of Posts and Telecommunications). Thanks for the hard efficient work of student Xudong Shi and the tutors Prof. Wendong Wang and Prof. Xiangyang Gong.

Valuable comment was received from Brian E Carpenter to improve the draft.

This document was prepared using 2-Word-v2.0.template.dot.

[Appendix A](#). Test Details of Host Behaviors




Figure 1 Test Environment

The 5 elements were all created in Vmware in one computer, for ease of operation.

- Router quagga 0.99-19 soft router installed on Ubuntu 11.04 virtual host
- DHCPv6 Server: dibbler-server installed on Ubuntu 11.04 virtual host
- Host A Window 7 Virtual Host
- Host B Ubuntu 12.10 Virtual Host
- Host C Mac OS X v10.7 Virtual Host

[A.1](#) Host Initialing Behavior

Host from non-configured to configured, we tested different A/M/O combinations in each OS platform. The states are enumerated as the following, 3 operation systems respectively:

- o Window 7
 - A=0&M=0&O=0, non-config
 - A=1&M=0&O=0, SLAAC only
 - A=1&M=0&O=1, SLAAC + Stateless DHCPv6
 - A=1&M=1&O=0, SLAAC + DHCPv6

- A=1&M=1&O=1, SLAAC + DHCPv6
- A=0&M=1&O=0, DHCPv6 only (A=0 or Non-PIO)
- A=0&M=1&O=1, DHCPv6 only (A=0 or Non-PIO)
- A=0&M=0&O=1, Stateless DHCPv6 only

o Linux/MAC OS X

- A=0&M=0&O=0, non-config
- A=1&M=0&O=0, SLAAC only
- A=1&M=0&O=1, SLAAC + Stateless DHCPv6
- A=1&M=1&O=0, SLAAC + DHCPv6
- A=1&M=1&O=1, SLAAC + DHCPv6
- A=0&M=1&O=0, DHCPv6 only (A=0 or Non-PIO)
- A=0&M=1&O=1, DHCPv6 only (A=0 or Non-PIO)
- A=0&M=0&O=1, non-config

As showed above, Linux and MAC OSX acted the same way, but differated from Windows 7. The only difference is when A=0&M=0&O=1, Windows 7 did stateless DHCPv6 while Linux/MAC OSX did nothing.

Result summary:

- A is interpreted as prescript in each OS at the initial state
- M is interpreted as prescript in each OS at the initial state
- O is interpreted as prescript in Windows 7
- A and M are independent in each OS at the initial state
- A and O are not totally independent in Linux and Mac, A=1 is required for O=1 triggering DHCPv6 info-request
- M and O are not totally independent in each OS. M=1 has the implication O=1

A.2 Host SLAAC/DHCPv6 Switching Behavior

- o SLAAC-only host receiving A=0&M=1
 - Window 7 would deprecate SLAAC and initiate DHCPv6
 - Linux/MAC would keep SLAAC and don't initiate DHCPv6 unless SLAAC is expired and no continuous RA
- o DHCPv6-only host receiving A=1&M=0
 - Window 7 would release DHCPv6 and do SLAAC
 - Linux/MAC would keep DHCPv6 and do SLAAC

When the host has been configured, either by SLAAC or DHCPv6, the operating systems interpreting the M flag quite differently. Windows 7 treats the flag as instruction, it even released DHCPv6 session when M=0. Linux and OS X were likely to treat the flag as advisory,

when SLAAC was done, it won't care about M=1, and M=0 won't cause operation for the already configured DHCPv6 addresses.

Please refer to [I-D.liu-6renum-dhcpv6-slaac-switching] for more details.

A.3 Host Stateful/Stateless DHCPv6 Behavior

- o StatelessDHCPv6-configured host receiving M=1 (while keeping O=1)
 - Window 7 would initiate stateful DHCPv6, configuring address as well as re-configuring other information
 - Linux/MAC no action
- o StatefulDHCPv6-configured host receiving M=0 (while keeping O=1)
 - Window 7 would release all DHCPv6 configurations including address and other information, and initiate stateless DHCPv6
 - Linux/MAC no action

Authors' Addresses

Bing Liu
Q14-4-A Building
Huawei Technologies Co., Ltd
Zhong-Guan-Cun Environment Protection Park, No.156 Beiqing Rd.
Hai-Dian District, Beijing
P.R. China

Email: leo.liubing@huawei.com

Ron Bonica
Juniper Networks
Sterling, Virginia 20164
USA

Email: rbonica@juniper.net