Network Working Group
Internet-Draft
Expires: August 23, 2008

B. Lim
C. Ng
K. Aso
Panasonic
February 20, 2008

Verification of Care-of Addresses in Multiple Bindings Registration
draft-lim-mext-multiple-coa-verify-01

Status of this Memo

This Internet-Draft will expire on August 23, 2008.

Copyright Notice

Abstract

Using multiple care-of address registration, there is a possibility that a malicious mobile node could create multiple care-of address bindings that does not belong to the mobile node at its home agent. The home agent would accept these bindings without verifying them due to the trust relationship it has with the mobile node. With these bindings, the mobile node can launch attacks by asking the home agent to flood the victims of these care-of addresses with useless packets. To mitigate such a problem, this memo introduces a verification mechanism that the home agent would use in order to verify the care-of addresses for the mobile node before using them for packet routing.

Table of Contents

1. Introduction

IP mobility allows a mobile node to maintain an ongoing communication session regardless of where the mobile node roams in the Internet. To achieve this constant connectivity, [1] permits the mobile node to register a transient IP address (e.g. care-of address) at an anchoring point (e.g. home agent). The mobile node sends a binding update (BU) message to the home agent for registering its care-of address. With this registration, the home agent could then forward any packets addressed to the mobile node at the mobile node's care-of address.

Security has always been a key consideration in user communications. This is especially so in IP mobility where the mobile node frequently changes its point of attachment. One such attack in IP mobility is the registration of an invalid care-of address at the home agent. Such registration allows the attacker to direct its traffic to other nodes. Such an attack is previously ignored since the attacker can only use one fake care-of address, thereby losing communications with its home agent once the fake address is registered. However, [2] permits the registration of multiple care-of addresses by means of using an identification number called Binding Unique Identification (BID) number to distinguish between multiple bindings to a home address. Unless proper security methodologies are implemented, the use of multiple care-of address registration heightens the chances of allowing an invalid care-of address to be registered at the home agent.

In Section 2, this document first expands on the problem briefly described in [2] and [3] about the security threat with the binding of fake care-of addresses. Next, possible approaches to mitigate such threat is analyzed in Section 3.

2. Problem Scope

For Mobile IPv6, it is assumed that the mobile node and its home agent establish some form of mutual authentication such that the home agent trusts the mobile node and will hence accept the registration of any care-of address that the mobile node presents to the home agent. This trust relationship is further strengthened when one assume that ingress filtering is being used such that when the home agent receives a binding update message from the mobile node stating its care-of address as the source address, the home agent trusts that the incoming packets do indeed originate from the specified source address. In addition, the home agent also trusts the routing infrastructure that packets forwarded by the home agent would be sent to the intended destination.

However, such a trust is no longer valid when the mobile node utilizes a single binding update message to register its multiple care-of address at the home agent. This technique is explained in [2] with the aim of introducing some optimization when registering multiple care-of address for a mobile node. Such optimization technique is useful in scenarios when resources (e.g. bandwidth) are scarce on some of the mobile node's interfaces, since it allows the mobile node to send a binding update message containing multiple care-of address to the home agent from an interface that does not have such resource constraint. Moreover, in Mobile IPv6, the use of the alternate care-of address option permits the mobile node to achieve the same effect of registering a care-of address for another interface via a specific interface. This introduces the risk of having fake care-of addresses registered at the home agent and compromise the security of the network. One applicable sceanrio that might encounter such a problem is the Long Term Evolution (LTE) system that is being worked on by the Third Generation Partnership Program (3GPP). Details on this scenario is shown in Appendix A.

There is a fundenmental difference between a mobile node using Mobile IPv6 to flood a victim with useless packets and a mobile node using Monami6 to flood a victim with useless packets. The difference lies in that using Monami6, the mobile node can bind a real care-of address at the home agent and use this care-of address to control the flow of attacks to the victims (e.g. to increase the packet transmission rate to the vicitm). In Mobile IPv6, the mobile node loses this means of control as the mobile node can only refresh the binding at the home agent. The example in Section 2.2 illustrates the esclated threat of Monami6 in more details.

2.1. Initiating the Attack

This problem is best understood through a specific example. Let us assume that a malicious mobile node, MN, sends a binding update message to its home agent, HA, to register multiple care-of addresses. This is illustrated in the Figure 1 below.

```
[Start of packet header]

Source Address           : CoA
Destination Address      : HA's address

[Mobility Options]

Binding Unique Identifier: BID1

Binding Unique Identifier: BID2
Care-of Address          : V1's address

Binding Unique Identifier: BID3
Care-of Address          : V2's address

[End of packet header]
```

Figure 1: Binding Update Message for Multiple Care-of Addresses Registration

CoA is a valid care-of address owned by MN. MN is attempting to bind addresses of two victims, V1 and V2, at HA in order to launch an attack towards the victims. The binding update message is secured using an IPsec security association to protect the integrity and authenticity.

When HA receives this binding update message, it will accept this binding update message based on the following. First, the binding update message is deemed authorized as the correct IPSec association key is used for the message. Second, the trust relationship that HA has with the routing infrastructure allows it to understand that this binding update message is sent from MN. Finally, after the first two checks have succeeded, the trust relationship that HA has with the MN permits it to trust the care-of addresses that are specified in this binding update message. Hence, the binding cache at HA will record three bindings for MN tied to MN's home address (HoA) as shown in Figure 2 below.

```
Binding 1 [HoA, CoA         , BID1]
Binding 2 [HoA, V1's address, BID2]
Binding 3 [HoA, V2's address, BID3]
```

Figure 2: Binding Cache at Home Agent

2.2. Launching the Attack

With the bindings created at HA, the malicious mobile node can now attempt to flood the victims with useless packets. This type of flooding can cause disruption of services at the victims' end as their devices will be busy processing these packets for meaningless data. This is illustrated in Figure 3.

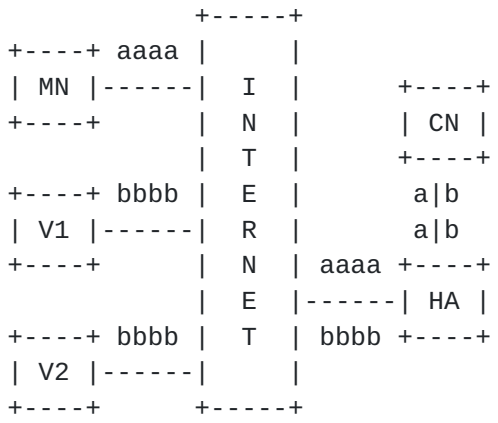

Figure 3: Packet flooding attack scenario

The binding cache at HA contains MN's care-of address along with V1's and V2's addresses. MN uses real-time transport protocol (RTP) and real-time transport control protocol (RTCP) to initiate a video stream from a streaming server (CN). Furthermore, MN could employ flow filtering techniques to assist in the forwarding of packets from HA. Such techniques are described in [4], [5] and [6].

Prior to starting the video flow, MN uses CoA to set filter rules at HA. One of the filter rule would tell HA that the control packets (shown as "a") for the video stream would be forwarded to BID1. Another filter rule informs HA that the data packets (shown in "b") for the video stream would be multicast to BID2 and BID3. When HA receives packets from CN, it would adhere to the filter rules created and forward the packets accordingly. This would cause V1 and V2 to be flooded with unintended data packets. Furthermore, MN can utilize care-of address1 to maintain the attack to V1 and V2 by sending binding update message to constantly renew the lifetime of the bindings. Likewise, MN can also use CoA1 to launch new attacks by

binding more victim's addresses at HA and modifying the filter rule to include them in the reception of the data packets.

Although both the home agent and the mobile node are mutually authenticated, this does not stop an authenticated mobile node from acting maliciously. Given time, the home agent might be able to detect that the mobile node is acting maliciously and may then deny mobility services to such a mobile node. One such way is when the home agent starts receiving requests from those victims to stop forwarding packets to it. With a substantial amount of such requests, the home agent might launch an investigation to determine if the mobile node was acting maliciously. Hence, we see that a home agent might not be oblivious to such an attack.

However, even if the home agent bars a single identity, the attacker may have multiple other identities available at his/her disposal. Take for example an electronic mail (e-mail) account, an attacker can create multiple e-mail accounts with same/different providers. Such creation can be done at any location that provides terminals with Internet connectivity (e.g. Internet cafe, Library) and does not require the attacker to log-in his/her actual identity. Using these multiple accounts, the attacker can start a spamming attack by flooding victims with useless e-mails. The attacker knows that eventually his/her account would get banned, but the attacker does not care. The purpose of the attacker has been achieved and that is to disrupt the services of those victims. Furthermore, the attacker can continue to create more accounts to replace those accounts that got banned. Thus, in this sense, having the identity known does not hinder the above problem from occurring.

3. Analysis: Possible Solution Approaches

Most of the soultions that will be described below assumes some form of ingress filtering is performed on the source address of the packet sent by the mobile node. Such assumption can be considered valid when we take in account that most operators deploying the 4G cellular networks would likely run ingress filtering within their networks as a security precaution to reduce the chances of packet spoofing.

One obvious way to approach the problem described in Section 2 is for the care-of addresses to be formed using the cryptographically generated addresses (CGA) technique [7]. With CGA in place, it permits the home agent to know if the mobile node actually owns that particular care-of address listed in the binding update message. Thus, the mobile node would be unable to bind fake care-of addresses at the home agent. However, introducing cryptographically generated care-of addresses increases the complexity of the mechanism to achieve multiple bindings. It is unclear how a message can contain multiple CGA signatures for each of the care-of addresses. Furthermore, complexity is increased by the fact that additional addresses are not found in the source address field but somewhere in the extension header of the packet (i.e. mobility header). This requires substantial integration between the CGA module and Mobility Support module in a network stack implementation. Additionally, each CGA Parameters structure which is at least 72 octets in length must be added to the binding update message, further increasing its size. Finally, the method of using CGA does not prevent the malicious mobile node to launch a flooding attack against a subnet by generating multiple non-existent care-of addresses in that subnet using the mobile node's own public keys.

Another approach to the problem, as described in [8], is to use an established symbiotic relationship between a mobile node and an access router in order to permit the home agent to verify the care-of address of the mobile node. The purpose of the symbiotic relationship between the access router and the mobile node allows the access router to act on behalf of the mobile node when it comes to validating the care-of address that the mobile node is using to the mobile node's home agent. The mobile node will tell its home agent of such symbiotic relationship establishment by providing the IP address of access router for the home agent to contact. This prompts the home agent to contact the access router to verify the care-of address that the mobile node claims to be using. Such message exchanges between the mobile node, home agent and access router are secured via public keys methodology which provides the proof to the home agent regarding the mobile node's care-of address. However, the requirement to include an access router role into such care-of address validation adds complexity to the access router. This is

especially so if the access router has a number of mobile nodes associated to it that requires such validation services. This implies that the access router would have to check and respond to the respective home agents on behalf of all these mobile nodes. Such action might significantly increase the processing load of the access router where such processing capabilities could be channeled towards more meaningful tasks such as packet routing, which is the primary purpose of an access router.

Another way to reduce the effect of the problem is to introduce some policy at the home agent to limit the amount of care-of addresses that a mobile node can bind. This limitation typically restricts a malcious mobile node from binding too many fake care-of addresses at the home agent, thus cuts down on the number of victims that a mobile node can affect. However, this approach does not stop the malcious mobile node from binding fake care-of addresses at the home agent. The intend of this approach is to control the after effect (e.g. mobile node triggering re-direction attack to vicitms), thereby isolating the problem to only a few victims rather than the whole system. Hence, this approach does not solve the problem.

Another alternative to solve the problem could employ the credit-based authorization (CBA) technique described in [9]. Here, the home agent uses the credits associated to an unverified binding to forward packets via that specific routing path. If that unverified binding does not have enough credits left, the home agent would not forward packets using that path as long as that binding remains unverified. The binding is verified when the home agent receives a packet from the mobile node using that care-of address as source. This technique limits the problem of packet flooding to a victim as it restricts the amount of data that can be transferred. For the home agent to fully utilize an unverified care-of address for packet routing the home agent would need to receive a packet from the unverified care-of address. With the reception of a packet from the unverified care-of address, this proves to the home agent that such a care-of address is addressable. However, the packets that the home agent receives from these unverified care-of address might not be sent by the mobile node. For example, it is possible that a malicious mobile node uses the Internet Control Message Protocol (ICMP) described in [10] to send a request message via its home address to unverified care-of address. If this unverified care-of address belongs to a victim, the reception of an ICMP request might trigger a response from the victim. Hence, the victim sends a response to the home address which, would be received at the home agent. With the reception of a packet from an unverified care-of address, the home agent would assume that the packet is sent by the mobile node. Thus, this verifies that care-of address for packet routing at the home agent. Also, in situations where an interface of the mobile node has

asymmetrical link characteristics (e.g. General Packet Radio System), that particular interface might have resource constraint on its upload path. Thus, such a constraint may make it costly for a mobile node to send a packet to its home agent using that particular interface.

Alternatively, the problem may be tackled by having the home agent send some encrypted information to the mobile node via the unverified care-of address and expects the mobile node to return the decrypted version of the information. This technique is described in [11] where the use of a cookie proves to the home agent that the mobile node is addressable at the specified care-of address. However, the optimization benefit of sending multiple care-of address within a single binding update is greatly reduced as it would require the mobile node to respond via all these care-of addresses. In this case, the mobile node might as well send the binding update message individually for each care-of address.

To optimize the inefficiency of using [11] for multiple care-of addresses case, it is possible for the home agent to send a notification to a mobile node via one unverified care-of address and ask the mobile node to respond to the reception of the notification via another care-of address. This utilizes the concept of multiple care-of addresses whereby the mobile node need not respond back using the same path as the reception of the request. Such a concept is particularly useful in the event that the mobile node has several unverified care-of addresses that needs to be tested. By asking the mobile node to respond via another unverified care-of address, in one round trip time, the home agent would be able to verify two care-of addresses. Figure 4 illustrates this.

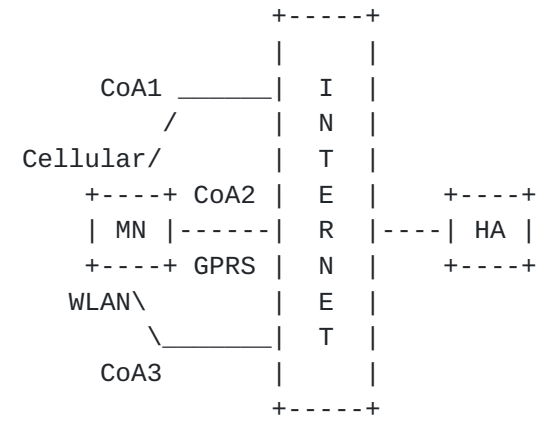

Figure 4: Operation of Home Agent

In Figure 4, the mobile node MN has three interfaces. A first interface associates with the cellular system and uses a care-of

address CoA1 for communication. A second interface connects to the General Packet Radio System (GPRS) and uses a care-of address CoA2 for communication. A third interface associates with the 802.11 Wireless Local Area Network (WLAN) and uses a care-of address CoA3 for communication. MN sends a binding update message from CoA1 that is attempting to register all its care-of addresses at its home agent, HA. Thus, when HA receives the binding update message, it binds CoA1 in its binding cache as a valid routing path to MN as it has proven to HA that MN is using that care-of address. However, for CoA2 and CoA3, HA would need to verify their addressability before using them for packet routing. Hence, HA uses the notification to verify both of these care-of addresses.

It is possible for the home agent to choose when to perform such verification process for the mobile node's bindings. One way is to have the home agent trigger the verification process immediately upon the reception of binding update message registering multiple care-of address from a mobile node. This is mostly helpful when it comes to supporting the concept of registering multiple care-of addresses. The purpose of having multiple care-of addresses is to allow the home agent to route packets to the mobile node via any of the multiple paths. Having the home agent perform the verification process immediately permits the home agent to quickly utilize all of the mobile node's care-of addresses for packet routing. For example, when the home agent receives the binding update message on CoA1 that specifies the binding of both CoA2 and CoA3, the home agent can choose to immediately verify these care-of addresses in order to start using them for packet routing to the mobile node.

Another way could be that the home agent triggers the verification process just before it needs to use an unverified routing path to the mobile node. This is especially so in the event that the mobile node sets a filter rule at the home agent specifying the routing of certain packets via the unverified routing path. When the home agent receives packets that match the filter rule condition, the home agent would trigger the verification process for the unverified path prior to using it. An example would be that the mobile node specify in a filter rule at the home agent that data packets would be routed via CoA3. When the home agent receives such a packet, it notes that CoA3 has yet been verified. Thus, the home agent performs the verification process to verify CoA3.

Also, the transmission path of the notification packet from the home agent can vary based on how strict the security policy is enforced. Strict security policy can forbid the home agent in using an unverified care-of address for any packet routing. In this sense, the home agent can only send the notification to the mobile node via a verified care-of address. Such policies are beneficial in

deterring any such flooding attacks being launch from the mobile node. For example, since CoA1 is the only verified care-of address to the mobile node, the home agent would send a message via CoA1 to notify the mobile node to respond back via CoA2. This action will cause the mobile node to understand that the home agent is trying to verify CoA2. Hence, the mobile node would send a packet to the home agent via CoA2, thereby proving to the home agent that it is addressable at CoA2.

On the other hand, if the home agent employs techniques such as CBA, it permits the home agent to route packets to the mobile node via an unverified care-of address. These packets would be limited to the amount of credits associated to that particular binding. The advantage of using an unverified path to transmit the notification is that it allows the home agent to verify at best two unverified care-of addresses. This is especially useful in the event that the mobile node has several unverified care-of addresses that needs to be tested. An example would be that the home agent sends a message via CoA2 to notify the mobile node to respond back via CoA3. This action allows the mobile node to understand that the home agent is trying to verify both CoA2 and CoA3. Hence, the mobile node would send a packet to the home agent via CoA3, thereby proving to the home agent that it is addressable at both CoA2 and CoA3.

Furthermore, it is possible that the home agent does not specify a return unverified care-of address when it notifies the mobile node. This would allow the mobile node to respond back to the home agent via any of the mobile node's care-of address. By permitting the mobile node to choose the respond path to the home agent, it supports the situation where the uplink path maybe resource constraint (e.g. limited uplink bandwidth). In the first place, this could be the reason why the mobile node decides to send a single binding update message to bind its multiple care-of address from a path that does not limit its resources. For example, the home agent wants to verify that CoA2 is addressable to the mobile node. Thus, the home agent notifies the mobile node via CoA2 to verify this care-of address. Since GPRS has limited uplink bandwidth, the mobile node decides to send the respond via the cellular path (CoA1). Hence, upon receiving the respond from the mobile node, the home agent is able to verify that the mobile node is addressable at CoA2.

Finally, how the notification information is carried to the mobile node may differ at the home agent. One possible way is that such notification is carried using a dedicated message format. This makes the notification packet small in size and easily recognized by the mobile node. On the other hand, since the notification is specifying just a care-of address that the mobile node should respond back by, an optimization to this would be to have the notification to be

tagged to packets that are being sent to the mobile node. One example could be the binding acknowledgment message that the home agent replies to the mobile node. Another example could be data packets that are addressed to the mobile node. By tagging it to existing packets, the home agent saves on the overhead incurred from the transmission the notification packet to the mobile node.

4. Conclusion

This draft explains that the use of multiple care-of address registration breaks the trust relationship between a home agent and a mobile node. With this relationship broken, it permits a malicious mobile node to bind multiple victims' care-of addresses at the home agent. With these bindings, the mobile node can launch attacks by flooding the victims with useless packets. To reduce the chances of having such a situation, we briefly analyse a few possible solutions that would be able to solve the problem. Base on the result of the analysis, we observed that a potential way to approach the problem in the multiple care-of address scenario is for the home agent send a message to the mobile node via one care-of address path and ask the mobile node to respond via another care-of address path. This technique not only allows the home agent to know that the mobile node owns the care-of address, but it further optimize the the amount of messages that the home agent needs to send in order to verify the mobile node's care-of addresses.

5. Security Considerations

This draft mentions new security requirements when implementing Multiple Care-of Address Registration [2].

6. IANA Considerations

TBD

7. References

7.1. Normative References

[1] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

[2] Wakikawa, R., "Multiple Care-of Addresses Registration", draft-ietf-monami6-multiplecoa-05 (work in progress), January 2008.

7.2. Informative References

[3] Montavont, N., "Analysis of Multihoming in Mobile IPv6", draft-ietf-monami6-mipv6-analysis-04 (work in progress), July 2007.

[4] Soliman, H., "Flow Bindings in Mobile IPv6 and Nemo Basic Support", draft-soliman-monami6-flow-binding-04 (work in progress), March 2007.

[5] Larsson, C., "A Filter Rule Mechanism for Multi-access Mobile IPv6", draft-larsson-monami6-filter-rules-02 (work in progress), March 2007.

[6] Mitsuya, K., "A Policy Data Set for Flow Distribution", draft-mitsuya-monami6-flow-distribution-policy-04 (work in progress), August 2007.

[7] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.

[8] Haddad, W., "Care-of Address Test for MIPv6 using a State Cookie", draft-haddad-mext-enhanced-reachability-test-00 (work in progress), February 2008.

[9] Arkko, J., Vogt, C., and W. Haddad, "Enhanced Route Optimization for Mobile IPv6", RFC 4866, May 2007.

[10] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[11] Dupont, F. and J. Combes, "Care-of Address Test for MIPv6 using a State Cookie", draft-dupont-mipv6-rrcookie-04 (work in progress), January 2007.

Appendix A. Applicable Scenario

The shift in the cellular system to utilize the Internet Protocol (IP) for communication and signaling implies the need for supporting IP mobility. For cellular operators, it means that a home agent would be located within their cellular system in order to support the mobile IP signaling (e.g. binding update) to their subscribers' mobile nodes. Within the cellular system, the establishment of a trust relationship between the operator and subscribers allows operators to provide services to subscribers. For example, the network would trust that the user operating a mobile node is the genuine subscriber and not someone who has stolen the subscriber's identity. To strength this trust relationship, operators uses the Authentication, Authorization, and Accounting (AAA) protocol within their system to authenticate their subscribers before providing services to them. Furthermore, a business contract (e.g. terms and condition) between the operator and subscribers binds subscribers to obey the rules dictated by the operator when using the services

provided by the operator. Using the terms and conditions, an operator is able to terminate services to a subscriber if the operator deems that the subscriber is acting maliciously. Figure 5 below illustrates the cellular scenario.

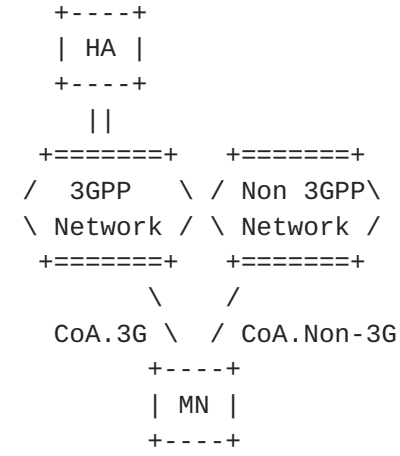

Figure 5: 3GPP scenario

In this scenario, a mobile node (MN) has two interfaces. A first interface is associated to a 3GPP network (e.g. 3G cellular) and uses CoA.3G for communication. A second interface is associated to a non-3GPP network (e.g. Wireless Local Area Network) and uses CoA.Non-3G for communication. The home agent in this scenario is part of the 3GPP network and might trust the validity of any packets sent with a source address configured in the 3GPP network. Thus, if MN sends a binding update message to HA with the source address as CoA.3G, HA would accept this binding as a valid route to the MN. However, this binding update message could further contain a request for HA to bind CoA.Non-3G.As HA trusts the packet with a source address CoA.3G, this implies that HA would trust its content, which in this case is the binding request for CoA.Non-3G. Hence, HA would bind CoA.Non-3G as a valid routing path to the MN even though the path has not been verified. This might incur the problem of a malicious MN using the HA to flood a victim with useless packets.

Similarly, in the event that the binding update message contains a mixture of care-of addresses that are trusted and not trusted by the home agent, the home agent could identify those trusted care-of addresses and omit performing the verification process for such addresses. When binding of a care-of address that is trusted by the home agent, the verification process could be skipped and the state of such a binding would immediately be set to the verified state. The reason that the home agent trusts this care-of address could be

that the home agent understands that the prefix used to create the care-of address is unique to the mobile node. Another reason could be that the care-of address is assigned by a trusted address allocation server (e.g. DHCP server). To be able to distinguish between trusted and non-trusted care-of addresses, policies could be installed at the home agent to assist in such identification process. If the home agent receives a binding update message that has a mixture of care-of addresses that are trusted and not trusted by the home agent, using the policy, the home agent would then be able to identify which care-of address requires verification and which does not. This reduces the chance of the home agent having to perform redundant care-of address verification on those trusted care-of address.

Appendix B. Change Log

- draft-lim-mext-multiple-coa-verify-01:
  - Reference lastest MCoA draft that specify the problem.
  - Added two more soultions (limiting care-of address binding at HA and CoA test using 'symbiotic' relationship) for analysis in Section 3.
- draft-lim-mext-multiple-coa-verify-00:
  - Initial version.

Authors' Addresses

Benjamin Lim
Panasonic Singapore Laboratories Pte Ltd
Block 1022 Tai Seng Ave #06-3530
Tai Seng Industrial Estate
Singapore 534415
SG

Phone: +65 65505478
Email: benjamin.limck@sg.panasonic.com

Chan-Wah Ng
Panasonic Singapore Laboratories Pte Ltd
Blk 1022 Tai Seng Ave #06-3530
Tai Seng Industrial Estate
Singapore 534415
SG

Phone: +65 65505420
Email: chanwah.ng@sg.panasonic.com

Keigo Aso
Matsushita Electric Industrial Co. Ltd. (Panasonic)
5-3 Hikarino-oka
Yokosuka, Kanagawa 239-0847
JP

Phone: +81 46 840 5123
Email: asou.keigo@jp.panasonic.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).